Adaptation aftereffects to facial expressions suppressed from visual awareness

Eunice Yang
Vanderbilt Vision Research Center and the
Department of Psychology, Vanderbilt University,
Nashville, TN, USA

Sang-Wook Hong
Vanderbilt Vision Research Center and the
Department of Psychology, Vanderbilt University,
Nashville, TN, USA

Randolph Blake
Vanderbilt Vision Research Center and the
Department of Psychology, Vanderbilt University,
Nashville, TN, USA, &
Department of Brain and Cognitive Sciences,
Seoul National University, Seoul, South Korea

The study of adaptation aftereffects has been used as a tool to investigate the neural events that give rise to face perception. Recent adaptation studies suggest that face processing does not occur outside of awareness since identity- and gender-specific face aftereffects cannot be induced when the adapting face is rendered perceptually invisible using interocular suppression. However there is substantial evidence suggesting that facial expression, unlike identity and gender, is an attribute of faces that may recruit processes that are engaged automatically and independent of observers' awareness and attention. Therefore we investigated whether adaptation aftereffects specific to facial expressions could arise under continuous flash suppression (CFS). Our results show that adaptation to facial expressions is virtually abolished, when faces are suppressed from awareness. Moreover, this loss in aftereffect strength cannot be attributed to contrast adaptation exclusively, since results show only modest changes in perceived contrast following face adaptation. When observers endogenously attend to the location of the suppressed adapting stimulus, expression-specific aftereffects are enhanced. Our findings suggest that neural activity specifying affective information of facial expressions is highly vulnerable to the disruptive effect of interocular suppression, but that allocation of attentional resources can partially counteract suppression's effect.

Keywords: facial expression, continuous flash suppression, aftereffects, emotion, awareness, attention

Citation: Yang, E., Hong, S.-W., & Blake R. (2010). Adaptation aftereffects to facial expressions suppressed from visual awareness. *Journal of Vision, 10*(12):24, 1–13, http://www.journalofvision.org/content/10/12/24, doi:10.1167/10.12.24.

Introduction

Faces provide a vital source of information for social communication and self-preservation. No wonder there is a wealth of evidence showing that face perception is a highly developed skill in primates (review by Tsao & Livingstone, 2008). Furthermore, face perception is supported by a complex neural network that readily encodes several face attributes (some dynamic), including identity, gender, race, and expression (Haxby, Hoffman, & Gobbini, 2000; Ishai, Schmidt, & Boesiger, 2005). The analysis of facial expressions, in particular, includes an additional affective component that may serve to prioritize the perception of certain objects or events (Anderson & Phelps, 2001; LeDoux, 1990).

Some have suggested that affective face processing may occur preattentively and outside of awareness (LeDoux, 1990; Morris, Ohman, & Dolan, 1998; Whalen et al., 1998 but see Palermo & Rhodes, 2007; Pessoa & Ungerleider, 2004). Several lines of evidence are consistent with this view, including results from studies employing techniques that induce interocular suppression (i.e., binocular rivalry and continuous flash suppression, or CFS: Tsuchiya & Koch, 2005). Specifically, several behavioral studies have shown that images of emotional faces, when in visual conflict with another stimulus viewed by the other eye, tend to achieve perceptual dominance rapidly (Jiang, Costello, & He, 2007; Yang, Zald, & Blake, 2007) and remain dominant for a greater percentage of the viewing period (e.g. Alpers & Gerdes, 2007; Yoon, Hong, Joormann, & Kang, 2009). Human brain imaging studies also provide evidence of neural activity in response to facial expressions that are rendered invisible with interocular suppression (Jiang & He, 2006; Pasley, Mayes, & Schultz, 2004; Williams, Morris, McGlone, Abbott, & Mattingley, 2004). A subcortical pathway has been posited as the neural circuitry underlying unconscious

doi: 10.1167/10.12.24 Received June 9, 2010; published October 21, 2010 ISSN 1534-7362 © ARVO

processing of affective cues (Morris, DeGelder, Weiskrantz, & Dolan, 2001), circuitry that may be specialized for rapid face encoding (Johnson, 2005). It is thought that this pathway may bypass the neural mechanisms underlying interocular suppression, particularly in early visual cortex (Lin & He, 2009) and is furthermore unsusceptible to the effects of visual attention (review by Johnson, 2005). However, the existence of such a pathway has yet to be confirmed anatomically.

If a subcortical pathway indeed mediates the analysis of emotional faces, the neural encoding of facial expressions may be impervious to the effects of interocular suppression and may transpire without awareness or attention. This claim would receive strong support if it were shown that the operations responsible for the neural representation of facial expressions adapted even under conditions of interocular suppression. One way to psychophysically measure neural adaptation under suppression is to examine visual aftereffects experienced following prolonged adaptation (Mollon, 1974; Thompson & Burr, 2009). We already know that adaptation to faces can induce reliable aftereffects of identity and gender. Specifically, a study by Moradi, Koch, and Shimojo (2005) had observers view faces under conditions that produced interocular suppression. Normally the identification of a test face is enhanced when it is presented subsequent to adaptation of its ‘anti-face’ in comparison to an unrelated face. However, this identity-selective aftereffect was abolished when the anti-face adapting stimulus was presented under suppression. Shin, Stolte, and Chong (2009) reported similar results for gender-specific adaptation aftereffects using another form of interocular suppression. Interestingly, when the authors directed observers’ spatial attention towards the location of the suppressed adapting face, the gender aftereffect occurred. Shin et al.’s results resonate with previous behavioral evidence of the attentional modulation of adaptation aftereffects under interocular suppression (Bahrami, Carmel, Walsh, Rees, & Lavie, 2008; Kanai, Tsuchiya, & Verstraten, 2006). Notably, both Moradi et al. and Shin et al. were able to reproduce low-level visual aftereffects under suppression, replicating previous studies (Blake, Tadin, Sobel, Raissian, & Chong, 2006). Thus the engagement of neural operations, such as those responsible for face identification and gender discrimination, is heavily dependent on observers’ perception of and attention to the face (Moradi et al., 2005; Shin et al., 2009). Given the evidence implicating a subcortical pathway for affective processing and the suggested pre-attentive characteristics of affective processing, it seems reasonable to ask how neural mechanisms responsible for the encoding of emotional facial expressions are impacted by interocular suppression.

Recently, Adams, Gray, Garner, and Graf (2010) adapted observers to emotional faces that were either visible or rendered perceptually invisible with CFS and then measured the facial expression aftereffect (FEA) induced by that period of adaptation. Interestingly, those authors found evidence of FEAs even though observers could not always perceive the adapting stimulus. The magnitude of FEAs produced from suppression however, were significantly smaller than those produced from the visible adaptation condition. Therefore, awareness modulated the strength of FEAs to some extent. While these findings are intriguing, trials in which the adapting face was visible and trials in which it was suppressed were presented randomly within the same session, making it possible that FEAs produced under CFS represented residual adaptation that carried over from visible adaptation periods. This seems highly possible given that observers experienced incomplete suppression on two thirds of suppression trials, thereby increasing the periods of visible adaptation.

Nonetheless, Adams et al.’s findings are consistent with the wealth of evidence implicating the distinct processing of affectively laden stimuli. We sought to replicate their findings in order to determine first, whether adaptation aftereffects to facial expressions could indeed be induced when observers were *completely* unaware of those expressions, owing to potent interocular suppression. Similar to Shin et al. (2009), we also investigated the effect of spatial attention on expression-specific aftereffects under interocular suppression. Moreover, we measured the aftereffect of adaptation to emotional faces on contrast perception, to determine whether a component of neural processing of emotional expressions during suppression are distinct from encoding of spatial contrast of the component features of those faces.

Experiment 1

Our first experiment asked whether adaptation to an image of an emotional face suppressed from awareness is able to influence subsequent judgments of emotional expression of test faces that are perceptually visible. The strength of this FEA was measured under conditions where the adapting face was always perceptually visible to observers and under conditions where the adapting face was continuously suppressed from awareness using CFS. Also included were conditions where the adapting stimulus was replaced with a blank field, to provide baseline measures against which to index FEA strength. The task for all conditions was the same: observers judged the expression of a test face chosen from a range of morphs between two facial expressions. The perceived categorical boundary between the two expressions (point of subject equality: PSE) was measured using an interleaved adaptive staircase procedure. Similar to previous studies (e.g. Fox & Barton, 2007; Hsu & Young, 2004), trials for the adaptation conditions and baseline condition were randomly intermixed within a session, which allowed observers to remain naïve of the condition being presented at any given time.

Method

Participants

Participants were drawn from the Vanderbilt University subject pool and from the local Nashville area. All were naïve to the purpose of the study and provided written consent. Eighteen participants with normal or corrected-to-normal acuity performed in Experiment 1. These participants were able to reliably suppress images according to our regime (below). An equal number of these participants (9) were assigned to one of two groups. Each group was presented with different faces and tested on a different continuum of facial expressions (i.e. happy–angry or angry–fearful).

Apparatus

Stimuli were presented on the left and right halves of a gamma-corrected CRT monitor (1024 × 768 resolution; 100 Hz frame rate) and viewed at a distance of 89 cm through a mirror stereoscope with mounted chin- and head- rests. Stimuli were controlled by a Macintosh G4 computer running MATLAB supplemented by the Psychophysics toolbox (Brainard, 1997; Pelli, 1997).

Stimuli

Stimuli consisted of eight faces (four female) taken from the Karolinska Database of Emotional Faces (Lundqvist & Litton, 1998). To confirm that the presence (or absence) of an FEA was not specific to a particular facial expression, we tested a different range of facial expressions for two groups of participants. For one group (Group HA), the adapting stimuli consisted of happy and angry expressions of four faces (2 female). Fearful and angry expressions of the remaining four faces (2 female) were used as adapting stimuli for the other group of participants (Group AF). Previous studies reported significant aftereffects using these pairs of facial expressions (e.g. Fox & Barton, 2007; Webster, Kaping, Mizokami, & Duhamel, 2004). Stimuli presented during the test period were morphs between the adapting facial expressions portrayed by the same individual. 100 morphs were created from each face at 1% intervals using Abrosoft Fantamorph 3 (www.fantamorph.com). Stimuli were cropped to remove features outside of the face (1.93° × 1.29°). All face images were scaled to gray and normalized to 20% contrast (root mean square). The CFS display consisted of colored Mondrian-like patterns (1.93° × 1.93°) that changed at a rate of 10 Hz. All stimuli were presented against a uniform background at mean luminance (33 cd/m^2).

Procedure

This experiment consisted of 2 sessions, with a given session testing three conditions concurrently: adaptation to 1) one facial expression in a pair (e.g. happy), 2) the opposing expression in the pair (e.g. angry), and 3) to a blank field (baseline condition). One session was devoted to trials where the adapting stimulus was always visible, and another session was devoted to trials where the adapting stimulus was always perceptually suppressed from awareness by CFS. The adapting stimulus, whether a face or a blank field, was always presented to the observer's sensory non-dominant eye, with the CFS display presented to the sensory dominant eye when suppression of the adapting face was required.

Regardless of session type, each session comprised a series of trials, each consisting of a brief adaptation period immediately followed by a test stimulus. A trial began with the presentation of binocular fusion contours and a fixation cross that served to promote stable binocular eye alignment. During trials of an adaptation condition, a face was presented to the center of the non-dominant eye upon key press. The contrast of the adapting face linearly increased during the initial 500 ms, in order to avoid abrupt transients. On baseline trials, no face stimulus was present during the adaptation period. In one session, a CFS display was presented to observers' dominant eye, which spatially and temporally overlapped with the presentation of the adapting face. In another session, there was no opposing stimulus in the dominant eye and the adapting face was clearly visible to observers.

The adaptation period lasted 5.5 s and was followed by a 50 ms blank interval. During the test period, a morph face was presented to both eyes for 200 ms at the visual field location of the previously presented adapting face. Observers were then given 3 s to categorize the expression of the morph using a two-alternative forced choice procedure (Figure 1). Observers initiated each trial with a key press. For each trial, the morph value (the proportion of each expression) was determined using a PEST adaptive staircase procedure (Taylor & Creelman, 1967). Each of the 3 conditions (2 adaptation and 1 baseline) was repeated four times, for a total of 12 staircases that were randomly interleaved within a session, with one constraint: the session involving CFS was always performed first to confirm that the observer was unable to perceive a face on those trials. During the adaptation period of a given trial in the CFS session, observers immediately pressed a designated key when they perceived anything other than the CFS display (they were not explicitly told that a face was present); this key press immediately terminated the trial and it was repeated when its staircase was selected next. If observers indicated disruption of suppression on more than 15% of trials, the experiment was terminated. Observers who completed the CFS session (70% of observers tested who had an average of 1% prematurely terminated trials) later confirmed that they had not perceived any faces during

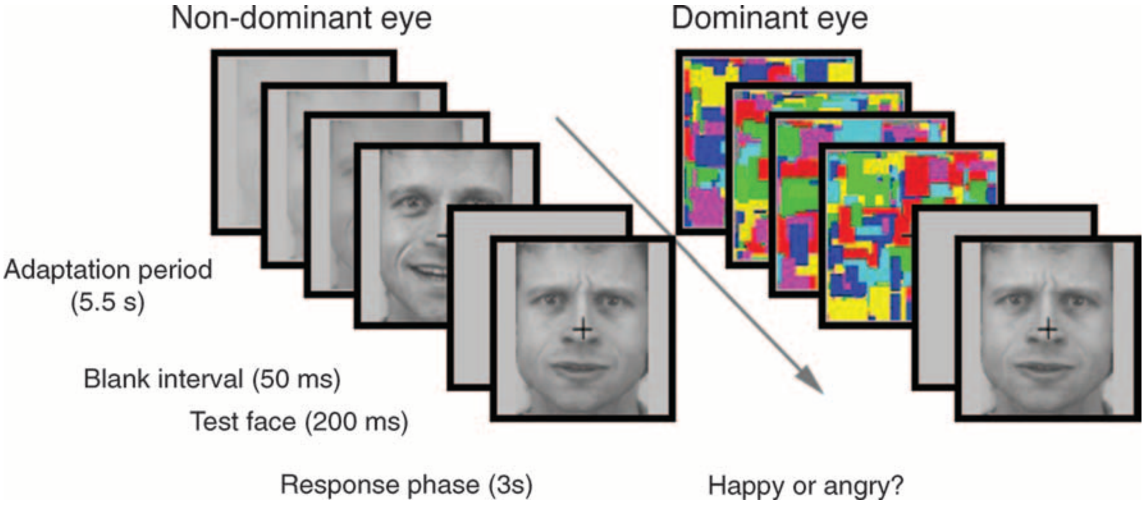

Figure 1. Example of an adaptation trial in Experiment 1 using CFS to perceptually suppress the adapting face (e.g. face with happy expression). After the adaptation period and a brief interval, a test face (e.g. morphed face of happy and angry expressions) was presented dioptically. Observers subsequently judged the expression of the test face by pressing one of two response keys. Trials of other conditions consisted of a blank frame in place of the CFS display (visible condition) or in place of the adapting face (baseline condition).

the CFS period upon questioning. They returned a week later to perform the same session but without CFS. Six practice trials were performed before each session, and a session took 60 min on average to complete.

Results: Experiment 1

FEA strength was expressed as the difference between PSE estimates measured on adaptation and baseline conditions. Of the 18 observers tested, a few failed to show an FEA when they adapted to a particular *visible* facial expression (happy: $n = 1$, fearful: $n = 2$, angry: $n = 3$). Their data for these particular facial expressions were excluded from further analysis. Results for Experiment 1 are summarized in Figure 2. Overall, there was a main effect of CFS on FEA magnitude ($F(1, 16) = 16.0$, $p = 0.001$) such that FEAs as a result of suppressed adaptation (mean strength ± standard error of the mean or *SEM*: 3.5 ± 1.2) were significantly weaker than FEAs measured under visible adaptation (10.8 ± 1.2). There was no session × group

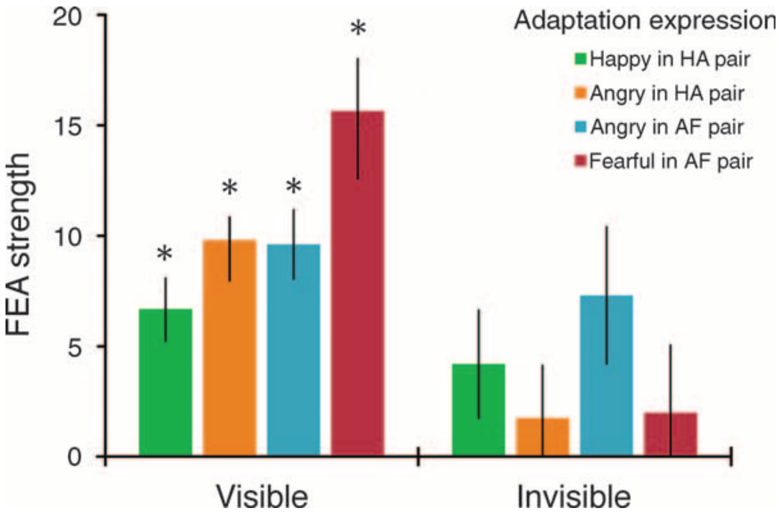

Figure 2. Results of Experiment 1: The magnitude of FEAs when adapting faces were visible or invisible due to CFS. Aftereffect strength was determined as the difference in PSE acquired from baseline (no adaptation) and from adaptation conditions. Values near 0 indicate no effects of face adaptation on PSE estimates of the test face. Asterisks indicate the adaptation conditions where FEAs were significantly greater than 0, after adjusting for multiple comparisons ($p < 0.006$). There were two groups of observers, one where happy and angry expressions served as adapting stimuli (HA pair) and the other where angry and fearful expressions were adapting stimuli (AF pair). Morphs between adapting expressions served as test faces. Error bars represent *SEM*. FEA = facial expression aftereffect.

interaction, indicating that this pattern of results occurred irrespective of which expression pair (happy–angry or angry–fearful) observers experienced during adaptation.

When the adapting stimulus was visible (no CFS), FEAs were significant in both Group HA (8.7 ± 1.2, $t(8) = 7.0$, $p < 0.001$) and Group AF (12.9 ± 2.1, $t(8) = 6.2$, $p < 0.001$). FEAs were observed when adapting to either expression in happy–angry pair (Happy: 6.7 ± 1.5, $t(7) = 4.5$, $p = 0.003$; Angry: 9.8 ± 1.9, $t(6) = 5.3$, $p = 0.002$) and in the angry–fearful pair (Angry: 9.6 ± 1.6, $t(7) = 6.0$, $p = 0.001$; Fear = 15.7 ± 3.1, $t(6) = 5.0$, $p = 0.002$). FEAs in the visible condition remained significant after correcting for multiple comparisons.

In contrast, when the adapting face was perceptually invisible, FEAs were not significantly different from 0 in Group HA. In Group AF, FEAs reached significance (4.5 ± 1.6, $t(8) = 2.8$, $p = 0.02$), specifically when the suppressed adapting face portrayed anger (7.3 ± 3.1, $t(7) = 2.3$, $p = 0.05$); however, it did not remain significant after adjusting for multiple comparisons, using Bonferroni correction. FEAs from the other adaptation conditions under CFS were not significant (HA—Happy: 5.6 ± 3.9; Angry: 3.4 ± 4.2; AF—Fear: 6.4 ± 8.0).

Experiment 2

Results from Experiment 1 indicate that FEAs are reduced to near-baseline levels when observers never perceived the adapting face owing to CFS. We wondered, however, whether this finding might be related to the generally weak magnitude of the FEAs measured in this experiment. Our experimental design involved randomly interleaving adaptation and baseline trials in a given session, a technique that could preclude strengthening of adaptation over time or, for that matter, carryover of adaptation into baseline trials. To avoid these possibilities, we performed a second experiment in which the adaptation and baseline conditions were administered in separate sessions, each beginning with an extended period of adaptation (to a face or a blank field). We also replaced the staircase procedure with the method of constant stimuli so that we could obtain an equal number of trials at all morph levels and, thus, generate complete psychometric functions for each condition.

Method

Participants

Eight observers participated in Experiment 2, none of whom participated in Experiment 1. The data from two observers was discarded since these observers were unable to reliably suppress the adapting stimuli, according to our criteria described earlier. All had normal or corrected-to-normal acuity. All participants provided written consent and with the exception one participant (author), all were naïve to the purpose of the study.

Apparatus and stimuli

The apparatus was the same as in Experiment 1. The happy–angry face pairs in Experiment 1 were used as adapting stimuli for Experiment 2. The corresponding test stimulus was a face morph that ranged between the full happy and angry expression (of the same person) at intervals of 5% (e.g. 35% Happy: 65% Angry), with a total of 21 possible morph values. The contrast, spatial dimensions and locations of the face and CFS stimuli were identical to those used in Experiment 1.

Procedure

Observers participated in 4 sessions, 2 devoted to the measurement of face judgments after adaptation (1 including CFS and the other without CFS) and 2 devoted to baseline measurements using the same test stimuli but adaptation to a blank field only (either accompanied by CFS or not). Two observers were exposed to happy facial expressions during adaptation periods and the remaining four observers adapted to the corresponding angry expressions of the same faces. Fusion contours and a fixation cross were always displayed to both eyes throughout a trial.

All sessions began with an initial, 2.4 min period of adaptation. If it was a face adaptation session, faces portraying an angry or happy expression were presented to observers' non-dominant eye in succession for 5.5 s (500 ms inter-stimulus interval). Each face identity (4) was repeated six times for a total 24 presentations of angry expressions. The contrast of each adapting face ramped up during the initial 500 ms to avoid abrupt transients. If it was a baseline adaptation session, a blank field within the fusion contour was presented to the non-dominant eye. For 2 sessions, a CFS display (10 Hz) was simultaneously presented to the dominant eye during this initial period (with or without an adapting stimulus in the non-dominant eye).

After this initial period, trials were identical to those presented in Experiment 1. Each trial began with 5.5 s of either top-up adaptation or a blank screen (with or without CFS in the opposing eye), followed by a 50 ms blank interval, and a 200 ms test morph stimulus. Observers had a 3 s time window in which to categorize the test stimulus as either happy or angry. Observers initiated each trial with a key press.

To find observers who experienced complete suppression of the adapting stimuli for the entire adaptation period (6 of 8 individuals tested), the first two sessions were always with CFS. In these sessions, observers were instructed to immediately indicate with key press whether they perceived anything other than the CFS display during the top-up adaptation period (without the assumption that faces in particular would appear). Those trials terminated and were presented again towards the end of the experiment (mean of 1% of total trials).

There were 20 trials for each morph value, yielding a total 420 trials per session. Test stimuli (morph value) and face identity were randomized and counterbalanced within a session. The order of adaptation and baseline sessions was also randomized for each observer, with the restriction that the first two sessions involved CFS. Each session was performed on consecutive days and took approximately 60 min per day to complete.

Results: Experiment 2

For each observer a best-fit Weibull function was first fit to his/her responses collapsed across the two baseline conditions. The category threshold was estimated from the morph value at which the two responses (happy and angry) were equally likely (PSE). Group psychometric functions were also generated from observers' data. In order to normalize and adjust for individual differences in baseline PSE values, individual psychometric functions were simply shifted such that each observer's baseline PSE fell on the midpoint of the morph values. Group PSE thresholds were then estimated for each adaptation condition. Data from the two baseline conditions were again combined since PSE values did not significantly differ (with CFS = 51.3; without CFS = 53.0), according to non-overlapping 95% bootstrap confidence intervals (see below). As in Experiment 1, FEA strength was determined by the difference in PSE thresholds between adaptation and baseline conditions (Figure 3 inset).

Figure 3 shows the group psychometric functions measured during adaptation and baseline sessions when CFS was present and when it was absent. When observers continuously perceived the angry expressions during adaptation, a significant FEA (15.6) was obtained as evidenced by the leftward shift of the psychometric function (red line) relative to the respective baseline (black line) condition. This means, in other words, that after adaptation observers tended to perceive the ambiguous, baseline morph as "happy" and thus required more of the angry expression in a morph to achieve the PSE. Similarly, a rightward shift in the psychometric function (blue line) occurred when observers adapted to happy expressions, with an FEA comparable in magnitude (13.6). The 95% confidence intervals derived from the bootstrapped threshold estimates (1000 repetitions; Efron & Tibshirani, 1993) confirm that, across as well as within observers, the thresholds from adaptation and baseline were significantly different (in addition, t-test on FEA scores: $t(5) = 11.9$, $p < 0.001$). Furthermore using the same stimuli, the FEA strength was significantly greater (14.9 ± 1.2, $t(13) = 3.3$, $p = 0.005$) and almost double in magnitude to that measured in Experiment 1 (8.7 ± 1.2), which suggests that the revised adaptation protocol in Experiment 2 was indeed more effective.

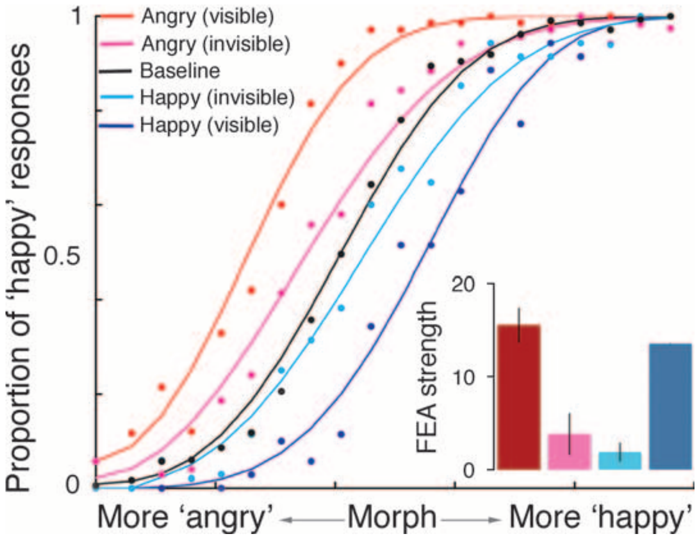

Figure 3. Results of Experiment 2: mean psychometric plots for 6 observers under conditions where CFS was present or absent during the adaptation period. Red and blue indicate conditions where an angry face or a happy face (respectively) was visible during the adaptation period. Pink and cyan indicate conditions where the adapting angry or happy face (respectively) was suppressed with CFS. Black is the mean performance for baseline conditions where no adapting face was present. Abscissa indicates the percentage of the happy expression in the morph test face where 0 is the angry expression. Ordinate displays the proportion of trials where observers' judged the test face as being happy. Inset shows the mean magnitude of FEAs across participants for each adaptation condition, relative to baseline.

Despite these larger FEAs associated with adaptation to visible faces, the magnitude of FEAs were not significantly different between Experiments 1 and 2 when adapting faces were invisible. In Experiment 2, FEAs measured after perceptual suppression of the adapting stimulus were again strongly attenuated across observers (angry condition: 4.5; happy condition: 4.9), as depicted by the pink and cyan lines in Figure 3. Consistent with this, FEAs from the invisible condition were significantly smaller than FEAs in the visible condition (11.7 ± 2.4, $t(5) = 4.9$, $p = 0.005$). Moreover, for no individual observer was the PSE value for the invisible condition significantly different from the PSE value for the baseline condition (overlapping 95% bootstrap intervals), implying that interocular suppression completely disrupted build-up of the FEA. When, however, we pooled results for all observers and applied a bootstrap confidence interval procedure, the two conditions were significantly different. It is fair to conclude, however, that this effect is marginal and, together with the results from Experiment 1, provide little evidence for residual adaptation during suppression.

Experiment 3

Results from Experiments 1 and 2 show that ordinarily robust FEAs are essentially abolished when the adapting stimulus is suppressed from awareness by CFS. This severe disruption in the build-up of the FEA was somewhat surprising to us, since several neuroimaging studies suggest that brain areas such as the amygdala respond to suppressed emotional faces as robustly as to visible emotional faces (Jiang & He, 2006; Pasley et al., 2004; Williams et al., 2004). This, in turn, led us to wonder whether the interocular suppression induced by our CFS display might be so potent that it erased essentially all traces of the face adaptation stimuli, including component features defined by luminance-defined contours. To evaluate this possibility, we turned to a different visual aftereffect that assesses adaptation's effect on perceived contrast of suprathreshold stimuli (Blakemore, Muncey, & Ridley, 1971), an aftereffect thought to tap into neural events transpiring relatively early in visual processing. We used the same adapting and test stimuli as those in Experiments 1 and 2, i.e., images of faces, and the design of Experiment 3 was similar to that of Experiment 2 in which each block of trials in the adaptation conditions began with an initial adaptation period followed by top-up adaptation. The major difference in this new experiment was the dependent variable being measured, i.e., apparent contrast.

Method

Participants

Seven observers (5 naïve about the purpose of the study) participated in Experiment 3. All had normal or corrected-to-normal acuity, and all provided written consent for participation.

Apparatus and stimuli

The apparatus was the same as that used in Experiments 1 and 2. Depending on observer, the face stimuli during the adaptation and test phases were either angry ($N = 4$) or happy ($N = 3$) facial expressions, which were the same stimuli as those used in Experiments 1 and 2. Stimulus parameters were also identical to those in previous experiments, with a few exceptions. Face and CFS stimuli were centered 1.29° to the left and/or right of fixation. The adapting face was presented at 20% RMS contrast. The corresponding pair of test faces was identical except that the contrast of one of the faces varied from trial to trial (probe), as determined by an adaptive staircase procedure (PEST). The remaining test face (pedestal) was always half of the adapting contrast. The face identity remained unchanged throughout a block of trials, and there were four different faces per expression (4 blocks per session).

Procedure

Observers participated in 3 sessions, 2 of which measured the contrast adaptation aftereffect and baseline performance (no adapting face) under CFS. These conditions were tested in separate blocks (2 per condition) that were counterbalanced and randomized per session. The third session measured the contrast adaptation aftereffect without CFS. Two sets of fusion contours (2.05° × 2.05°) were presented 1.29° to the left and right of fixation dioptically during stimulus presentations.

Blocks of trials involved an initial adaptation period of 60 s, followed by trials consisting of 5.5 s of top-up adaptation, a 300 ms ISI and finally, a 200 ms test phase. Depending on the condition, a face or a blank frame (baseline) was presented during the initial and top-up adaptation periods. When an adapting face was present, it was randomly presented within one of two fusion frames and to observers' non-dominant eye for each block. The contrast of the adapting face ramped up (500 ms) during its initial presentation to avoid abrupt transients. In the first two sessions, CFS displays were presented to both hemifields in observers' dominant eye during the initial and top-up adaptation periods. In the final session, the adapting stimulus was presented without a competing stimulus in the dominant eye. Following the blank interval, faces during test were presented in both hemifields to the adapted eye. On half of the staircases for a given condition, the probe face was presented in the same location as the adapting face and the pedestal face was presented in the opposite hemifield. On the remaining staircases, the locations of the probe face and pedestal face were switched. Observers were given 3 s to indicate which one of the two faces was higher in contrast via key press. Subsequent trials were initiated with another key press. There was a mandatory 90 s break between blocks.

Each block consisted of 2 staircases (same condition and identity) that were randomly interleaved, and a staircase terminated once the PSE between the test and reference stimulus contrast was reached. Since there were 4 blocks per session and 3 sessions total, 8 threshold values were measured for each of the 3 conditions. As in Experiments 1 and 2, observers terminated a trial if the adapting face broke suppression, and that trial was later repeated (mean was 1% of total trials). The 3 sessions were performed on separate days and took 45 min on average to complete.

Results: Experiment 3

Using the same face adaptation stimuli as those used in our two previous experiments, Experiment 3 examined whether CFS weakens the induction of an aftereffect that presumably involves neural adaptation of local, contrast-defined features of the face. As in Experiments 1 and 2, the magnitude of the contrast adaptation aftereffect was derived from the difference in matching contrast values under each adaptation condition relative to the baseline condition.

There was a significant change in apparent contrast relative to baseline when observers were aware of the adapting face (22.9 ± 1.7%, $t(6) = 13.3$, $p < 0.001$) and, of particular interest, when the adapting face was suppressed from awareness completely (15.1 ± 1.2%, $t(6) = 12.8$, $p < 0.001$; Figure 4). The aftereffect measured following suppression of the adapting face was significantly weaker than the effect produced when the adapting face was visible ($t(6) = 3.3$, $p = 0.02$). There were no significant differences between observers who were presented with an angry adapting face and those who were presented with a happy adapting face.

Results from Experiment 3 show that the test face presented at the location of the adapting face was perceived as weaker in contrast relative to the test face presented at an adjacent location where the adapting face was not present. This result was obtained whether or not observers perceived the adapting stimulus, implying that significant contrast adaptation occurs even when the adapting face was suppressed from awareness. This finding is reminiscent of binocular rivalry suppression's weakening influence on the threshold elevation aftereffect measured using grating patterns (Blake et al., 2006). The present results imply that the virtual abolishment of FEAs in Experiments 1 and 2 reflects the attenuation of neural responses beyond those involved in registering contrast-defined feature components of a face, such as the orientation and curvature of the contours defining the eyes and the mouth. Note, we are not claiming that contrast adaptation plays no role in mediating the FEA. In fact, face aftereffects are most strongly induced when visual features of the adapting and test face are held constant (e.g. Ellamil, Susskind, & Anderson, 2008; Fox & Barton, 2007), which suggests that the adaptation of low-level features may indeed be one component of the induction of face aftereffects. Rather, we are suggesting that contrast adaptation cannot account entirely for the magnitude of weakening of the FEA when the adapting stimulus is suppressed. Evidently, neural mechanisms beyond those involved in contrast perception are engaged when we view an emotional face, and activity in those mechanisms are substantially attenuated when a face is suppressed by CFS.

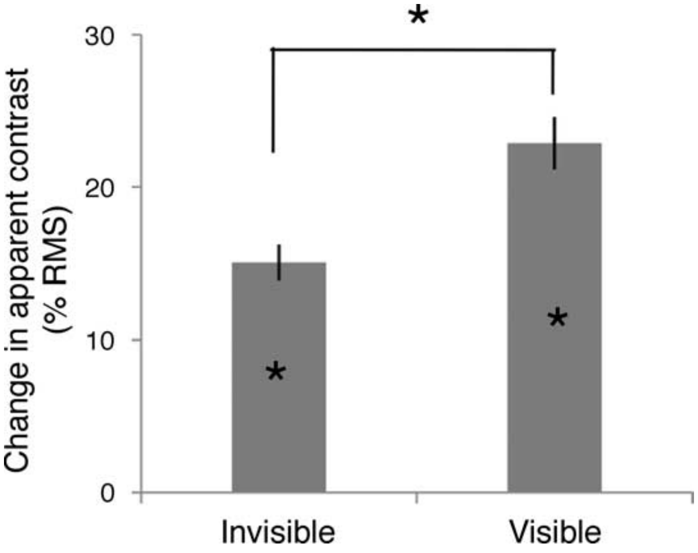

Figure 4. Results of Experiment 3. Change in apparent contrast (% RMS) of the test face relative to baseline when the adapting face was visible or rendered invisible by CFS. Error bars represent *SEM*. Asterisks within bars indicate conditions significantly differently from 0 and asterisks between bars indicate significant differences between conditions.

Experiment 4

Recent studies have shown that attention can modulate the strength of aftereffects produced by adaptation to stimuli falling outside of awareness (Bahrami et al., 2008; Kanai et al., 2006; Shin et al., 2009). In particular, Shin et al. (2009) reported that the gender-adaptation aftereffect did not occur unless observers' spatial attention was directed towards the suppressed face during CFS. Those findings motivated us to perform a final experiment to determine whether FEAs could be similarly strengthened under conditions where observers' spatial attention is directed toward the location of a suppressed adapting stimulus.

Method

Participants

Eight observers (2 authors and 6 naive) participated in Experiment 4. Observers had normal or corrected-to-normal acuity and provided written consent.

Apparatus and stimuli

The same apparatus was used as in previous experiments. Face and CFS stimuli were also the same as those used in Experiments 1–3. Specifically, adapting stimuli were faces with happy expressions and test stimuli were faces morphed along the continuum of happy–angry expressions of each person (1 male, 1 female). The stimulus display was identical to Experiment 3. Two sets of fusion contours (2.05° × 2.05°) were presented 1.29° to the left and right of fixation and on both halves of the monitor during the entire session. Stimuli were presented within these contours during trials. The attention task was adapted from the study by Shin et al. (2009).

Procedure

A session began with an initial adaptation phase of 2.2 min. During the first 300 ms, a green line (with one end on fixation) pointed either to the left or right of fixation. This line served as a cue for observers to covertly attend to one of two locations that were centered at 1.29° to the left and right of fixation. An adapting stimulus (face with happy expression) was then presented to one of these locations to the non-dominant eye for 5.5 s (ramped up in contrast over 500 ms) across all trials in the initial adaptation period. The location of the adapting stimulus (left versus right) remained the same across the entire experiment and was randomly chosen for each observer. The cued location overlapped with the location of the adapting stimulus 50% of the time. At the same time, a CFS display (10 Hz) was simultaneously presented to both locations (left and right of fixation) to the dominant eye. The CFS sequence that was consistent with the direction of the cue contained one or two frames (minimum ISI of 800 ms, after initial 600 ms of CFS) that were on average, 30% lower in luminance relative to other CFS frames. During this adaptation period (5.5 s), observers' were instructed to attend covertly to the cued CFS display while maintaining fixation and to count the number of dimmed frames appearing during the CFS sequence (discrimination task). Once the stimuli were removed from the screen, observers had 1 s to respond by pressing the '1' or '2' key (with their left hand) before the next adaptation trial began. When the observers' response was correct on a given trial, a tone was presented after their key press.

Following the initial adaptation phase, randomized trials of the adaptation and baseline conditions were administered (Figure 5A). Each trial began with a cue pointing to the location either left or right of fixation for 300 ms. During adaptation trials, the adapting stimulus was presented in the same location across trials (consistent with initial adaptation phase, e.g. left) to the non-dominant eye for 5.5 s and the cue pointed to this location. During trials of

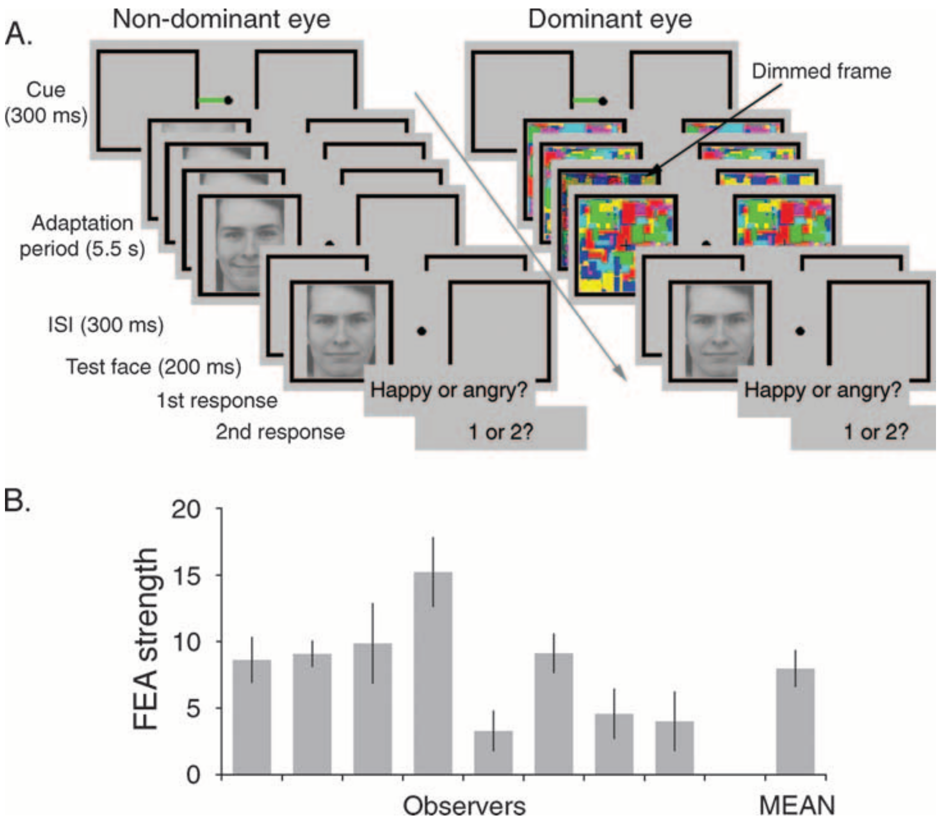

Figure 5. (A) Example of an adaptation trial in Experiment 4. Observers were cued to covertly attend to the left or right of fixation. During top-up adaptation, a happy face was presented in the cued location to the non-dominant eye while CFS images were presented in both locations to the dominant eye. A test face was briefly presented in the cued location and observers were to first respond to the expression of the test face and subsequently report the number of dimmed CFS frames (1 or 2) that had occurred during the adaptation period. (B) Experiment 4 results. FEA strength is plotted relative to baseline for each observer when performing an attention task in the location of a suppressed adapting face. First two data points on the left represent authors' data. Error bars represent *SEM*.

the baseline condition, no adapting stimulus was presented during this period and the cue pointed to the location opposite to do the adapting face in the adaptation trials (e.g. right). In both conditions, CFS displays were simultaneously presented in both locations to the dominant eye. Identical to the initial adaptation phase, observers performed a discrimination task in which they counted the number of dimmed frames (one or two) in the cued CFS display while maintaining fixation. Following a 300 ms ISI, a test face was dioptically presented in the cued location for 200 ms. This location was the same as the adapting stimulus during trials of the adaptation condition (e.g. left) and for baseline trials, the location was opposite of the face stimuli in the adaptation condition (e.g. right). A response screen with "Happy or Angry" was then presented at fixation and observers pressed the left or right arrow key with their right hand to indicate the perceived expression in the test face. Once observers responded to the test face, another response screen with "1 or 2" was presented. Observers indicated the number of dimmed CFS frames using the '1' or '2' key with their left hand and feedback was given. The next trial immediately began after both responses were chosen.

For each trial, the morph value (the proportion of each expression) was determined using PEST. There were for 4 staircases for each stimulus identity (2) and therefore 8 staircases per condition (2). All staircases were randomly interleaved and each terminated once the PSE was reached. As in previous experiments, observers were instructed to terminate a trial if the adapting face broke suppression (mean of 2.7% of total trials). Staircases were separated into 2 sessions (8 per session; 4 per condition). Observers took approximately 60 min per session to complete the experiment.

Results: Experiment 4

Mean performance on the CFS discrimination task was 77% correct, confirming that observers' performance did not reach ceiling and the task required close attention to perform. In Figure 5B, FEA strength was calculated as the difference in thresholds (PSEs) acquired from adaptation and baseline conditions. Values greater than 0 indicate that FEAs were observed with our manipulation of observers' spatial attention.

Indeed, relative to baseline PSE (66.2 ± 2.9), observers' PSE were significantly lower (i.e. required more of the angry expression in the ambiguous, baseline morph) when they attended to the location of a suppressed happy face (58.2 ± 3.7, $t(7) = 5.7$, $p = .001$). When comparing FEAs observed in this experiment with FEAs produced by the same stimuli in previous experiments, FEAs here were significantly greater than those from Experiment 1 ($t(15) = 2.4$, $p = .03$) and from Experiment 2 ($t(12) = 2.3$, $p = .04$). However unlike Experiment 2, we were able to find significant FEAs with CFS on an individual basis, for 7 out of 8 observers ($ps < .05$). These results suggest that FEAs are reliably strengthened if observers' spatial attention is oriented toward the location of the invisible adapting face.

Discussion

Our findings reveal that facial expression aftereffects are strongly modulated by observers' awareness and attention during face adaptation, providing little evidence for the involvement of preattentive, subcortical processing of affective visual stimuli in the task used in this study. FEAs were virtually abolished when the adapting face stimulus was perceptually suppressed using CFS (Experiments 1 and 2). Contrast adaptation aftereffects induced by the same suppressed faces were observed and the attenuation in aftereffect strength was insufficient to account completely for the effect of CFS on FEAs. This pattern of results suggests that CFS disrupts processing in neural mechanisms including those uniquely related to processing of affective facial information (Experiment 3). On the other hand, we found that covertly directing observers' spatial attention towards the location of the suppressed adapting stimulus facilitated face adaptation, thereby inducing larger FEAs with CFS (Experiment 4).

Our results add to the wealth of evidence showing weakened visual aftereffects caused by interocular suppression of a visual adaptation stimulus (e.g. Blake et al., 2006; Moradi et al., 2005; Shin et al., 2009; Tsuchiya & Koch, 2005). The tendency in those studies, incidentally, is for the reduction in aftereffect strength to vary depending on the complexity of the adapting stimulus (e.g., gratings vs. faces) and the nature of the task used to assess visual aftereffects (e.g., detection vs. identification). Of particular relevance, our results are consistent with findings of abolished identity- and gender-specific face aftereffects under interocular suppression (Moradi et al., 2005; Shin et al., 2009). In contrast, Adams et al. (2010) observed FEAs when observers were unaware of the adapting face stimuli, although they were attenuated relative to FEAs produced from visible adaptation. In their study, on a majority of 'suppression' trials, observers reported intermittent periods at which the adapting face was visible, indicating that suppression during adaption was incomplete. In our study, only trials of complete suppression were examined and those occurred nearly all the time for observers. Furthermore with the exception of Experiment 2, baseline trials were always intermixed with adaptation trials in order to minimize demand characteristics. Their version of interocular suppression is more comparable to binocular rivalry in which the dominance and suppression of a given stimulus perceptually alternate with time. FEAs may have been induced solely by periods where the adapting face was visible in dominant and suppression trials. This adaptation carry-over effect may also explain why authors found significant FEAs with briefly visible periods (500 ms) of the adapting face since these trials were intermingled with trials of prolonged visible adaptation (4 s) in their control experiment. In our study, visible and suppression trials were presented in separate sessions to avoid adaptation carry-over effects. Alternatively, given that rivalry suppression reduces contrast sensitivity less so than does CFS (Tsuchiya, Koch, Gilroy, & Blake, 2006), it is possible that face stimuli used by Adams et al. (2010) were sufficiently high in contrast to partially overcome the effect of suppression.

Previous neuroimaging studies have reported robust neural activity in response to suppressed facial expressions, particularly in the face-selective areas of the superior temporal sulcus (Jiang & He, 2006) and the amygdala (Jiang & He, 2006; Pasley et al., 2004; Williams et al., 2004), areas associated with affective processing. On the other hand, our results are consistent with reported neural behavior in another face-selective area, the fusiform face area (FFA), where the activity observed during suppression is noticeably reduced in magnitude relative to activity when faces were perceptually visible (Jiang & He, 2006; Sterzer, Haynes, & Rees, 2008). Similarly, attenuated stimulus-specific responses during interocular suppression are also observed in areas relatively early in the visual processing stream, including the primary visual cortex and the thalamus (review by Rees, 2007). Therefore, our results seem to dovetail with the pattern of neural activity and adaptation aftereffects implicated along the ventral visual cortical pathway rather than within the subcortical pathway.

Our findings also suggest that the neural operations responsible for the FEAs induced by suppressed adapting faces, whether cortical or subcortical, may be strongly modulated by attention. While the long-standing assumption is that affective processing transpires without the engagement of attentional mechanisms (e.g. Vuilleumier, Armony, Driver, & Dolan, 2001), emerging evidence suggests that the processing of affective cues may be modulated by attention. Expression-selective responses may only arise when attentional resources are available to process those faces. For instance, when observers' attention was consumed by another task, neural activity to different facial expressions could not be dissociated in face-selective areas including the 'fear'-sensitive amygdala (Pessoa & Ungerleider, 2004). Similarly, the typical attenuation of neural activity in response to repeated faces (repetition suppression) in fMRI-adaptation paradigms only occurs when those faces are attended to (e.g. Henson & Mouchlianitis, 2007). Kouider, Eger, Dolan, and Henson (2009) reported similar attentional modulations when attended faces were made invisible with backward masking. Consistent with this, studies using interocular suppression show that attentional resources must be available and purposely directed to the suppressed adapting stimulus in order for aftereffects to transpire (Bahrami et al., 2008; Kanai et al., 2006; Shin et al., 2009). Altogether these findings suggest attentional resources must be available and further allocated to the location of the face stimulus for adaptation aftereffects to occur, even when the face is outside observers' awareness. Our results are consistent with this interpretation.

Our findings imply that endogenous attention directed to the location of an otherwise invisible face may successfully boost neural signals (e.g. contrast gain), facilitating neural activity associated with the adaptation producing FEAs. Indeed, there is evidence that covert attention enhances neural activity in response to face stimuli (e.g. Vuilleumier et al., 2001; Williams, McGlone, Abbott, & Mattingley, 2005; Wojciulik, Kanwisher, & Driver, 1998). There is some indication that areas become increasingly contrast invariant along the visual processing hierarchy, with face- and object-selective areas displaying similar response magnitudes to low- versus high-contrast images (e.g. Avidan et al., 2002). Interestingly, Murray and He (2006) found that this contrast invariance in object-selective areas was modulated by endogenous attention; when observers' spatial attention was directed away from an object, the pattern of activity in these higher-order areas varied as a function of stimulus contrast and was similar to the pattern observed in V1. We surmise that our manipulation of spatial attention (Experiment 4) may have effectively boosted signals in early (e.g. V1) and subsequent stages (e.g. FFA) of visual analysis, to levels sufficient to produce strong, reliable face-adaptation aftereffects. We believe our findings have now set the stage for future studies and one question worth examining is whether this attention-modulated aftereffect is non-retinotopic.

In conclusion, our findings suggest that the processes dedicated to the analysis of affective information conveyed by facial expression are susceptible to the effects of interocular suppression, just like those involved in the analysis of face identity and gender. Furthermore, the strength of residual signals associated with a suppressed face stimulus can be boosted by spatial attention directed to the location of that stimulus.

Acknowledgments

The authors would like to thank Michael Bailey for his assistance in testing participants. This work was supported by NIH grants EY13358 and 5T32 EY007135 and by the WCU program through the National Research Foundation of Korea funded by the Ministry of Education, Science and Technology (R32-10142).

Commercial relationships: none.

Corresponding authors: Eunice Yang and Randolph Blake.
Email: eunice.yang@vanderbilt.edu; randolphblake@gmail.com.
Address: Department of Psychology, Vanderbilt University, PMB 407817, 2301 Vanderbilt Place, Nashville, TN 37240-7817, USA.

References

Adams, W. J., Gray, K. L. H., Garner, M., & Graf, E. W. (2010). High-level face adaptation without awareness. *Psychological Science, 21,* 205–210.

Alpers, G. W., & Gerdes, A. B. M. (2007). Here is looking at you: Emotional faces predominate in binocular rivalry. *Emotion, 7,* 495–506.

Anderson, A. K., & Phelps, E. A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature, 411,* 305–309.

Avidan, G., Harel, M., Hendler, T., Ben-Bashat, D., Zohary, E., & Malach, R. (2002). Contrast sensitivity in human visual areas and its relationship to object recognition. *Journal of Neurophysiology, 87,* 3102–3116.

Bahrami, B., Carmel, D., Walsh, V., Rees, G., & Lavie, N. (2008). Spatial attention can modulate unconscious orientation processing. *Perception, 37,* 1520–1528.

Blake, R., Tadin, D., Sobel, K. V., Raissian, T. A., & Chong, S. C. (2006). Strength of early visual adaptation depends on visual awareness. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 4783–4788.

Blakemore, C., Muncey, J. P., & Ridley, R. M. (1971). Perceptual fading of a stabilized cortical image. *Nature, 233,* 204–205.

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

Efron, B., & Tibshirani, R. (1993). *An introduction to the bootstrap*. London: Chapman & Hall.

Ellamil, M., Susskind, J. M., & Anderson, A. K. (2008). Examinations of identity invariance in facial expression adaptation. *Cognitive, Affective & Behavioral Neuroscience, 8,* 273–281.

Fox, C. J., & Barton, J. J. S. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research, 1127,* 80–89.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4,* 223–233.

Henson, R. N., & Mouchlianitis, E. (2007). Effect of spatial attention on stimulus-specific haemodynamic repetition effects. *Neuroimage, 35,* 1317–1329.

Hsu, S. M., & Young, A. W. (2004). Adaptation effects in facial expression recognition. *Visual Cognition, 11,* 871–899.

Ishai, A., Schmidt, C. F., & Boesiger, P. (2005). Face perception is mediated by a distributed cortical network. *Brain Research Bulletin, 67,* 87–93.

Jiang, Y., Costello, P., & He, S. (2007). Processing of invisible stimuli: Advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological Science: A Journal of the American Psychological Society/APS, 18,* 349–355.

Jiang, Y., & He, S. (2006). Cortical responses to invisible faces: Dissociating subsystems for facial-information processing. *Current Biology, 16,* 2023–2029.

Johnson, M. H. (2005). Subcortical face processing. *Nature Reviews Neuroscience, 6,* 766–74.

Kanai, R., Tsuchiya, N., & Verstraten, F. A. J. (2006). The scope and limits of top-down attention in unconscious visual processing. *Current Biology, 16,* 2332–2336.

Kouider, S., Eger, E., Dolan, R., & Henson, R. N. (2009). Activity in face-responsive brain regions is modulated by invisible, attended faces: Evidence from masked priming. *Cerebral Cortex, 19,* 13–23.

LeDoux, J. (1990). *The emotional brain*. New York: Simon & Schuster.

Lin, Z., & He, S. (2009). Seeing the invisible: The scope and limits of unconscious processing in binocular rivalry. *Progress in Neurobiology, 87,* 195–211.

Lundqvist, D., & Litton, J. (1998). The Averaged Karolinska Directed Emotional Faces—AKDEF, CD ROM from Department of Clinical Neuroscience, Psychology Section, Karolinska Institutet.

Mollon, J. (1974). After-effects and the brain. *New Scientist, 61,* 479–482.

Moradi, F., Koch, C., & Shimojo, S. (2005). Face adaptation depends on seeing the face. *Neuron, 45,* 169–175.

Morris, J. S., DeGelder, B., Weiskrantz, L., & Dolan, R. J. (2001). Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain: A Journal of Neurology, 124,* 1241–1252.

Morris, J. S., Ohman, A., & Dolan, R. J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature, 393,* 467–470.

Murray, S. O., & He, S. (2006). Contrast invariance in the human lateral occipital complex depends on attention. *Current Biology, 16,* 606–611.

Palermo, R., & Rhodes, G. (2007). Are you always on my mind? A review of how face perception and attention interact. *Neuropsychologia, 45,* 75–92.

Pasley, B. N., Mayes, L. C., & Schultz, R. T. (2004). Subcortical discrimination of unperceived objects during binocular rivalry. *Neuron, 42,* 163–172.

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

Pessoa, L., & Ungerleider, L. G. (2004). Neuroimaging studies of attention and the processing of emotion-laden stimuli. *Progress in Brain Research, 144,* 171–182.

Rees, G. (2007). Neural correlates of the contents of visual awareness in humans. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 362,* 877–886.

Shin, K., Stolte, M., & Chong, S. C. (2009). The effect of spatial attention on invisible stimuli. *Attention, Perception & Psychophysics, 71,* 1507–1513.

Sterzer, P., Haynes, J., & Rees, G. (2008). Fine-scale activity patterns in high-level visual areas encode the category of invisible objects. *Journal of Vision, 8*(15):10, 1–12, http://www.journalofvision.org/content/8/15/10, doi:10.1167/8.15.10. [PubMed] [Article]

Taylor, M., & Creelman, C. (1967). PEST: Efficient estimates on probability functions. *Journal of the Acoustical Society of America, 42,* 1097.

Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology, 19,* R11–R14.

Tsao, D. Y., & Livingstone, M. S. (2008). Mechanisms of face perception. *Annual Review of Neuroscience, 31,* 411–437.

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience, 8,* 1096–1101.

Tsuchiya, N., Koch, C., Gilroy, L. A., & Blake, R. (2006). Depth of interocular suppression associated with continuous flash suppression, flash suppression, and binocular rivalry. *Journal of Vision, 6*(10):6, 1068–1078, http://www.journalofvision.org/content/6/10/6, doi:10.1167/6.10.6. [PubMed] [Article]

Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2001). Effects of attention and emotion on face processing in the human brain: An event-related fMRI study. *Neuron, 30,* 829–841.

Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature, 428,* 557–561.

Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B., & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience, 18,* 411–418.

Williams, M. A., McGlone, F., Abbott, D. F., & Mattingley, J. B. (2005). Differential amygdala responses to happy and fearful facial expressions depend on selective attention. *Neuroimage, 24,* 417–425.

Williams, M. A., Morris, A. P., McGlone, F., Abbott, D. F., & Mattingley, J. B. (2004). Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. *Journal of Neuroscience, 24,* 2898–2904.

Wojciulik, E., Kanwisher, N., & Driver, J. (1998). Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. *Journal of Neurophysiology, 79,* 1574–1578.

Yang, E., Zald, D. H., & Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion, 7,* 882–886.

Yoon, L., Hong, S. W., Joormann, J., & Kang, P. (2009). Perception of facial expression of emotion during binocular rivalry. *Emotion, 9,* 172–182.