

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

A novel face aftereffect based on recognition contrast thresholds

İpek Oruç*, Jason J.S. Barton

Human Vision and Eye Movement Laboratory, Departments of Medicine (Neurology) and Ophthalmology and Visual Sciences, University of British Columbia, Vancouver, BC, Canada

ARTICLE INFO

Article history:

Received 11 January 2010

Received in revised form 26 April 2010

Keywords:

Face recognition

Adaptation

Aftereffects

Contrast thresholds

Response suppression

Sharpening

ABSTRACT

Previously, repulsive perceptual-shift face aftereffects have been reported. Here, we introduce a novel face adaptation method involving changes in contrast thresholds for face recognition. We find non-monotonic changes for adapted faces, with facilitation at short and suppression at long durations. Thresholds for unadapted faces were unaffected at short but elevated at long durations, more than those for the adapted face. A population-coding model showed that selective suppression of adapted representations cannot explain repulsive perceptual-bias aftereffects. The findings indicate greater complexity to adaptation, with facilitation, suppression, lateral inhibition of unadapted representations, and additional perceptual factors at long durations.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

Adaptation aftereffects are changes in perception induced by a preceding stimulus. Aftereffects are widespread in the visual system, occurring for both low-level properties such as luminance, contrast, spatial frequency, orientation, and motion (Anstis, Verstraten, & Mather, 1998; Blakemore & Campbell, 1969; Gibson & Radner, 1937), and for higher-order representations such as shapes and faces (Clifford & Rhodes, 2005; Leopold, O'Toole, Vetter, & Blanz, 2001; Suzuki, 2005). Experimentally, aftereffects are revealed by at least two phenomena (e.g., Blakemore & Nachmias, 1971). The first is a perceptual shift, usually a 'repulsive' aftereffect, in which the perception of a subsequent stimulus is shifted away from the properties of the preceding adapting stimulus. For example, in orientation, adaptation to a stimulus with counter-clockwise tilt will cause a subsequent vertical stimulus to appear tilted clockwise (Gibson & Radner, 1937). The second is a change in detection thresholds: typically, viewing the adapting stimulus causes stimuli with properties similar to the adaptor to become harder to perceive (Blakemore & Campbell, 1969; Regan & Beverley, 1985).

To explain both types of aftereffects, perceptual shifts and threshold changes, many use models in which the percept reflects the net population response of a large number of individual units, with each unit responsive to only a limited range of values of the

stimulus property in question (e.g., orientation). In such models, adaptation is often explained by response suppression. That is, during sustained viewing of an adapting stimulus with a particular value for that property (e.g., 10° counter-clockwise tilt), the responses of units preferring that value are reduced, impairing the detection of subsequent stimuli with similar values, while units responding to other values (e.g., 45° clockwise tilt) are unaffected. In addition to the predictable effects on thresholds, a second result of this selective reduction in response is that the net population response for the next stimulus (e.g., vertical tilt) is shifted away from suppressed values, resulting in a 'repulsive' perceptual shift (e.g., towards clockwise tilt) (Clifford, Wenderoth, & Spehar, 2000; Coltheart, 1971; Mather, 1980).

Recently repulsive perceptual-shift aftereffects have also been demonstrated for faces. Adapting to one facial identity biases the perception of a subsequent face away from this identity (Fox, Oruc, & Barton, 2008; Leopold, O'Toole, Vetter, & Blanz, 2001). Similar aftereffects have been shown for facial properties such as expression, gender, ethnicity (Fox & Barton, 2007; Webster, Kaping, Mizokami, & Duhamel, 2004), viewpoint (Fang & He, 2005), and gaze direction (Jenkins, Beaver, & Calder, 2006). These face aftereffects are not due to aftereffects for lower-level image properties such as contrast, size, or tilt (Butler, Oruc, Fox, & Barton, 2008), as they persist despite changes in image size (Zhao & Chubb, 2001), retinal location (Fang & He, 2005; Leopold, O'Toole, Vetter, & Blanz, 2001) and viewpoint (Jiang, Blanz, & O'Toole, 2007).

Because perceptual-shift aftereffects derive from the change in the 'relative' balance of activity in units that respond preferentially to the adapting stimuli versus those that do not, they cannot inform us of the 'absolute' changes in the responses of these units

* Corresponding author. Address: Human Vision and Eye Movement Laboratory, VGH Eye Care Center, 2550 Willow Street, Vancouver, BC, Canada V5Z 3N9. Fax: +1 604 875 4302.

E-mail address: ipor@interchange.ubc.ca (İ. Oruç).

induced by adaptation. If viewing of a specific face makes the observer less likely to perceive the properties of this face in a subsequent ambiguous test face, is this due to suppression of the representations of the adapting face, facilitation of the representations of unadapted faces, or both? To address this question experimentally, the second type of aftereffect – changes in luminance contrast threshold – may be useful. In this technique, the independent variable is not the amount of face A versus face B mixed in the ambiguous test stimulus, but the amount of luminance contrast, which is therefore orthogonal to the relationship between A and B.

In this study, our first goal was to use contrast thresholds for face recognition to determine how adaptation affects recognition performance. Classical models of adaptation based on selective response suppression would predict that thresholds for recognizing the adapted face should be elevated, while those for recognizing unadapted faces should not be affected (Coltheart, 1971; Graham, 1989; Mather, 1980). The contrast threshold technique provides an instrument to directly test this account.

Our second goal was to use this technique to determine the temporal dynamics of adaptation effects. Typically, lower-level aftereffects (e.g., contrast, tilt, or visual motion) grow monotonically with increasing adapting duration in the range of seconds to minutes, whether they are measured as perceptual shifts (Magnussen & Johnsen, 1986) or elevated detection thresholds (Blakemore & Campbell, 1969). Results for faces from two studies so far, both using perceptual-shift aftereffects, also show a logarithmic increase in aftereffect magnitude when adapting duration was increased from 1 s to 16 s (Leopold, Rhodes, Muller, & Jeffery, 2005; Rhodes, Jeffery, Clifford, & Leopold, 2007). However, very brief adapting durations (<500 ms) were not used, although both neurophysiological recordings (Albrecht, Geisler, Frazor, & Crane, 2002) and psychophysical studies (Suzuki, 2005) have shown substantial effects for non-face stimuli with brief adaptation, e.g., improved orientation discrimination following adaptation for periods of 400–500 ms (Dragoi, Sharma, Miller, & Sur, 2002; Muller, Metha, Krauskopf, & Lennie, 1999). Furthermore, there is growing evidence that higher-level aftereffects may emerge at shorter adapting durations (Fang, Murray, Kersten, & He, 2005; Kohn, 2007; Suzuki, 2005) than their lower-level counterparts. Thus, in the present study we measured aftereffects following adapting durations of 10–6400 ms to extend our knowledge of the temporal dynamics of face aftereffects, and to determine this for both adapted and unadapted faces.

2. Experiment 1: a contrast-based face aftereffect

2.1. Methods

2.1.1. Observers

Seven observers with normal or corrected-to-normal vision participated (2 males, ages 25–35), of which four participated in the main experiment (1 male, ages 25–32), and four participated in the control experiment (1 male, ages 28–35). With the exception of IO, who participated in both experiments, all observers were naïve to the purposes of the experiment. The protocol was approved by the review boards of the University of British Columbia and Vancouver Hospital, and informed consent was obtained in accordance with the principles in the Declaration of Helsinki.

2.1.2. Apparatus

Stimuli were displayed on a SONY Trinitron 17-inch GDM-G500 monitor at 1024 × 768 resolution and refresh rate of 100 Hz. Viewing distance was 99 cm. Cambridge Research Systems (CRS) VSG Toolbox for Matlab was used to present the stimuli via a CRS VSG 2/3 card. Displays were gamma-corrected by means of an

automated calibration procedure using the VSG software and an OptiCAL photometer (Model OP200-E) by CRS. Average luminance was 40 cd/m².

2.1.3. Stimuli

Face stimuli were five female faces displaying a neutral expression obtained from the Karolinska Database of Emotional Faces (Lundqvist & Litton, 1998). All face images were first converted to grayscale using Adobe Photoshop CS 8.0 (www.adobe.com). The images were then manipulated using Matlab (www.math-works.com). An oval aperture was superimposed on the face images, outside of which the display was uniform gray set at the mean luminance of 40 cd/m². The tip of the nose and the pupils were used as anchors to align faces horizontally and vertically. All faces had the same pose (frontal), tilt (vertical), and eye-color (brown), without obvious distinguishing marks such as moles and visible hair, to minimize discrimination based on trivial features. Luminance values inside the oval aperture were normalized such that the average was set to the mean luminance and the root-mean-squared (rms) contrast (the standard deviation of luminance values divided by mean luminance) to one. These images were the standard templates whose contrasts were later modified by the staircase procedure. The horizontal and vertical extents of the cropped faces were 5.1 × 7.2° visual angle, respectively, at the viewing distance of 99 cm.

2.1.4. Procedure

Face recognition contrast thresholds at the 82% correct level were estimated using a five-alternative forced-choice (5-AFC) paradigm. In each trial, one of the five alternative faces was selected randomly as the test stimulus and shown for 150 ms. The observer was then required to indicate which of the five faces the test stimulus resembled. A psychophysical staircase implemented in Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) based on the Quest procedure (Watson & Pelli, 1983), controlled the contrast of the test face at each trial. An estimate of the threshold is obtained at the end of a fixed 40-trial run per staircase.

Each trial started with an adaptation period during which either one of the five possible faces at 60% contrast or a blank stimulus was shown. The adapting duration was fixed within a block of trials, but differed across blocks. Adapting durations used were 10 ms, 20 ms, 50 ms, 100 ms, 200 ms, 400 ms, 800 ms, 1600 ms, 3200 ms and 6400 ms. The order of blocks (i.e., adapting durations) was randomized for each subject. At each trial the adaptation period was followed by a white noise mask (50 ms), a fixation cross (150 ms), a blank screen (150 ms), a test stimulus (150 ms), a blank screen (150 ms), and finally the choices screen, which remained visible until the subject responded with a keypress (see Fig. 1). Auditory feedback indicated whether the response was correct. A new trial started as soon as the observer made their keypress. There were 30 possible adapting-test pairs (6-adapting stimuli × 5-test stimuli). A separate contrast threshold was measured for each adapting-test pair, by using 30 randomly interleaved staircases, each controlling one adapting-test pair. In addition, filler trials that contained any one of the 30 adapting-test pairs, with a test stimulus at very low contrast, were randomly interspersed throughout the block with 1/6 probability, to prevent the observer from forming strategies based on tracking the progress of the 30 experimental staircases. The responses to filler trials were discarded.

The procedure for the control experiment performed to exclude contributions from low-level retinotopic properties was identical except that the adapting face was 50% larger than the test faces (i.e., 7.7 × 10.8°), and the test faces were presented 1° left or right of central fixation, determined randomly at each trial, so that the




Fig. 1. Illustration of a typical experimental trial. Each trial started with an adapting period, varying between 10 ms and 6400 ms in different blocks. There were six possible adapting stimuli: one of five faces or a blank stimulus. The adapting period was followed by a 50 ms white noise mask, a 150 ms fixation cross, and a 150 ms blank. This was followed by the test stimulus (150 ms), which was randomly chosen from the same set of five faces. With six adapting stimuli and five test stimuli, there were 30 possible adapting stimulus-test pairs. The test stimulus was followed by a 150 ms blank screen that was replaced by an answer display. The answer display remained on the screen until the observer responded.

corresponding facial regions did not fall upon the same portions of the visual field.

The observers were instructed to avoid fixating at a set point on the face and were encouraged to scan the face to prevent negative afterimages to form, especially at the longer adapting durations. The entire experiment lasted about 12–14 h for each subject, and was divided into 30–60 min-long sessions on different days. Due to the lengthy nature of the blocks with the longer adapting durations, some sessions ended mid-staircase, and were continued into the next session. All sessions started with a 10-trial warm-up block which featured the same stimuli, experimental procedure, and adapting duration to allow the subjects settle into the experimental environment and task before starting the experimental block. Responses during the warm-up block were discarded and were not included in the analysis. Prior to participating in the experiment, all subjects were familiarized with the task and the stimuli by completing at least 5000 trials in a 5-AFC discrimination task using the same face stimulus set.

2.1.5. Data analysis

The 30 adapting-test pairs were categorized into three conditions: *congruent*, where the adapting and test faces were the same; *incongruent*, where the adapting and test faces were different; and *baseline*, where the adapting stimulus was a blank. A *threshold ratio* was calculated by dividing the threshold for an adapting condition by the threshold for the baseline condition. Thresholds at the baseline condition did not vary across adapting durations ($F(9, 27) = 0.98, p = 0.48$), therefore a baseline value averaged from all adapting durations was used. Threshold ratios for the *congruent* and *incongruent* conditions were computed by taking geometric means

of all ratios obtained in that condition: these ratios were calculated for each adapting duration. If adaptation has no effect on the thresholds, then the threshold ratio is expected to be one. An elevation (i.e., suppression) is indicated by threshold ratios larger than one, and a decrease (i.e., facilitation) is indicated by threshold ratios between 0 and 1. Group data were computed by taking geometric means of threshold ratios across subjects at each adapting duration. Non-parametric bootstrap confidence intervals (Efron & Tibshirani, 1993) were computed by re-sampling the data a large number of times with replacement. The resulting sample is then ordered and the 95% confidence intervals are given by the lower and higher 2.5th percentile values.

2.2. Results

2.2.1. Main experiment

The effect of adaptation on the recognition of the same face seen during adaptation (*congruent condition*) was complex and non-monotonic (Fig. 2). Compared to a baseline 'no-adaptation' condition, a brief adaptation of at least 20 ms facilitated recognition performance, in that it decreased recognition contrast thresholds, whereas at durations greater than 400 ms adaptation suppressed recognition, in that it increased contrast thresholds. For the *incongruent condition* (test face different from the adapting face), adaptation had no effect on thresholds at very short durations of 10–20 ms, but with longer periods it generated increases in thresholds that grew as adapting duration lengthened. A comparison of *congruent* and *incongruent* conditions showed that, after at least 20 ms of adaptation, thresholds were always higher in the *incongruent condition*, until the thresholds for both *congruent* and *incon-*




Fig. 2. Results. Threshold ratio is plotted as a function of adaptation duration. Group data are computed as geometric means of threshold ratios across all observers. Black and red line colors indicate the congruent and incongruent conditions, respectively. The dashed line indicates the baseline of one, at which there is no effect of adaptation. Error bars indicate 68% bootstrap confidence intervals. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

gruent conditions converged at an adapting period of about 6400 ms.

The data for adaptation durations around 5000 ms are relevant to the reports of perceptual-shift aftereffects, which most frequently use adaptation periods of 5000 ms (Butler et al., 2008; Fang & He, 2005; Fox & Barton, 2007; Jiang, Blanz, & O'Toole, 2006; Rhodes & Jeffery, 2006). Explanations of perceptual shifts based upon selective response reduction suppose that adapted representations are suppressed and unadapted representations are unaffected by having viewed the adapting face. However, our results clearly show that at this duration of adaptation, recognition contrast thresholds for both adapted and unadapted faces are elevated, indicating suppression for both types of representations. Furthermore, if anything, the thresholds for unadapted faces are more rather than less elevated than those for adapted faces.

2.2.2. Control experiment

To determine if these effects are due to adaptation of low-level image properties such as local orientation and contrast, we performed a second control experiment. Most aftereffects for low-level properties are retinotopic: that is, they are localized to the region of the visual field where the image of the stimulus was projected, and require steady fixation during the adaptation period to emerge. Our subjects were allowed and even encouraged to shift gaze during the adaptation period, minimizing the contribution of retinotopic low-level aftereffects; however, for the trials with brief viewing durations the conditions may not have permitted more than a single fixation. Other strategies to exclude low-level effects have been to show that perceptual-shift aftereffects persist despite changes in the retinal location (Fang & He, 2005; Leopold et al., 2001) or image size (Zhao & Chubb, 2001) between adapting and test stimuli, both of which alter the retinotopic correspondence of the two images.

Our control experiment used both strategies. Adapting faces were 50% larger in size than the test faces, and viewed centrally, while test faces were the same size as in the main experiment, but centered 1° to the left or right of central fixation, determined randomly in each trial. We collected control data at two adapting durations: 50 ms and 1600 ms. These durations are the points at which the main characteristics of the temporal profile can be observed, namely the early facilitation, and later suppression for

the congruent curve, and the monotonic increase in thresholds with increasing adapting duration for the incongruent curve.

The results of the control experiment (Fig. 3) show that the overall temporal pattern of the aftereffect is preserved in the control experiment. In the congruent condition there was significant facilitation after 50 ms adaptation (Wilcoxon signed-rank test, $p < 0.05$), which reversed at the long adapting duration (1600 ms) into a significant elevation of threshold (Wilcoxon signed-rank test, $p < 0.05$), a pattern consistent with the results of the main experiment. The difference between short (50 ms) and long (1600 ms) adapting durations was also significant, as revealed by a pair-wise comparison between the two ($p < 0.05$, Wilcoxon rank-sum test).

The pattern for the incongruent condition seen in the first experiment was also reproduced. There was no significant adaptation effect at the short adapting duration ($p > 0.1$), but a significant increase in thresholds at the long adapting duration ($p < 0.05$, Wilcoxon signed-rank test). Pair-wise comparison between the short and the long adaptation also showed a trend towards an increase in threshold elevation at the long adapting duration ($p = 0.08$, Wilcoxon rank-sum test).

Pair-wise comparisons between the corresponding conditions in the control and main experiments reveal no significant differences between the two (Wilcoxon rank-sum test, p 's > 0.1) except for a slight reduction in effect size for the incongruent condition at 1600 ms adapting duration ($p < 0.05$). Overall, then, the characteristics of the aftereffect temporal profile are preserved in the control experiment. Therefore these aftereffects are not likely due to low-level retinotopic adaptation. Consistent with previous studies of face aftereffects (Butler et al., 2008; Fang & He, 2005; Leopold et al., 2001; Zhao & Chubb, 2001), we conclude that face adaptation involves representations beyond local retinotopic distributions of contrast, orientation, or other low-level image properties.

2.3. Comment

We explored a large range of adapting durations to determine the temporal profile of aftereffects for both adapted and unadapted faces. We found aftereffects with adapting durations as short as 20 ms. The temporal profile for adapted faces was non-monotonic, showing threshold reduction (facilitation) at short adapting durations, and threshold elevation (suppression) at longer durations,

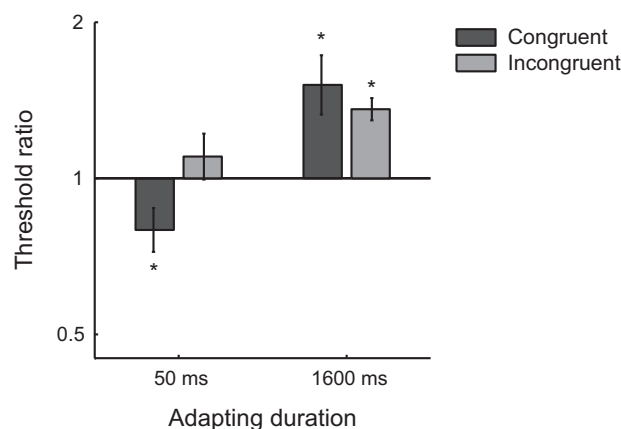


Fig. 3. Results of the control experiment. Threshold ratios at short (50 ms) and long (1600 ms) adapting durations are shown. The data represent geometric means of threshold ratios across four observers. Error bars indicate 68% bootstrap confidence intervals. Data for congruent and incongruent conditions are shown in dark-gray and light-gray bars, respectively. Conditions that differ significantly from the baseline value of one in the expected direction are marked with an asterisk (all p 's < 0.05).

the latter consistent with hypotheses of response suppression for adapted representations. However, we found threshold elevations for unadapted faces at long durations also. This differs from classical models of adaptation that postulate that suppression is specific to adapted stimuli (Coltheart, 1971; Graham, 1989; Mather, 1980). Indeed, in low-level vision, empirical findings show that aftereffects disappear when adapting and test stimuli are too dissimilar (e.g., Blakemore & Campbell, 1969; Regan & Beverley, 1985). The reasons why faces do not behave in a similar fashion in our data are not yet clear. One possibility is that, since faces are highly similar stimuli, unadapted faces share in the aftereffects of the adapted faces. However, if this were the case, one would expect that the aftereffects in the *incongruent condition* should be qualitatively similar to but quantitatively smaller than those in the *congruent condition*. That is, the *incongruent condition* should show a modest facilitation at short durations and a modest elevation of thresholds at long durations, compared to the *congruent condition*. Instead, we find no facilitation at short durations and, if anything, a larger elevation of thresholds at long durations.

Other models offer alternative explanations. In their priming model, Huber and O'Reilly (2003) postulated that contributions from stimulus persistence, accommodation, and lateral inhibition might bring about a complex relationship of adaptation effects with adaptation duration (Fig. 4A). Their model proposes that short exposure to an adapting stimulus would facilitate performance for a similar test stimulus, due to persistence of the representation of

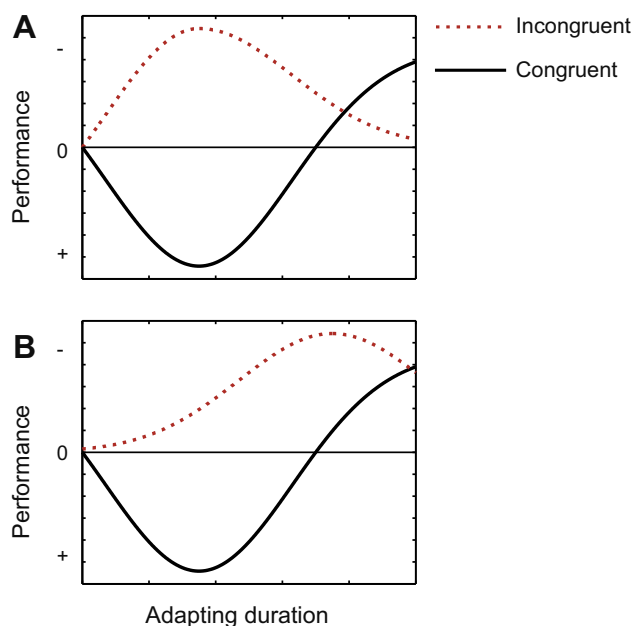


Fig. 4. Model results based on Huber and O'Reilly (2003). (A) Performance change with respect to adapting duration is plotted based on the model of Huber and O'Reilly (2003). The solid black line represents the *congruent condition*, and the red dotted line represents the *incongruent condition*. According to Huber and O'Reilly (2003), brief pre-exposure to a similar stimulus (solid black line, corresponding to our *congruent condition*) facilitates performance due to stimulus persistence. With continuing pre-exposure, accommodation reverses this effect and results in diminished performance at longer durations. For the *incongruent condition*, the opposite pattern is predicted at brief pre-exposure due to lateral inhibition. As the facilitation at the *congruent condition* (i.e., adapted face) diminishes at longer durations, so does the inhibitory effect, and thus the performance returns to baseline levels for the *incongruent condition*. (B) If lateral inhibition takes time (500–1000 ms) to develop, as suggested in studies of other stimuli (Andrews, 1967; Carpenter & Blakemore, 1973), this would shift the predicted temporal profile for the *incongruent condition* forward along the adapting duration axis. The pattern of temporal dynamics predicted in this case resembles our experimental results (Fig. 2). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the adapting stimulus. At longer durations this effect reverses into suppression for similar stimuli, due to accommodation. Our results for the adapted face are consistent with these predictions.

On the other hand, their model predicts that facilitation for adapted stimuli should be mirrored by increased thresholds for unadapted stimuli, due to lateral inhibition from the facilitated stimulus. Lateral inhibition should resolve as facilitation declines and is replaced by accommodation, at which point aftereffects for unadapted stimuli should be negligible. Instead, we find only a minimal increase in thresholds for unadapted stimuli at durations when facilitation is greatest for adapted stimuli, and substantially elevated thresholds at durations where facilitation has been replaced by accommodation for adapted stimuli. One possible explanation of this discrepancy is a temporal lag between the dynamics of persistence and lateral inhibition, so that the maximal effect of the latter is delayed relative to the former. Indeed, lateral inhibition in the orientation domain is known to take 500–1000 ms to develop fully (Andrews, 1967; Carpenter & Blakemore, 1973). If there is a similar delay in the build-up of competitive interactions between adapted and unadapted face representations, the threshold curves for unadapted faces would be shifted forward along the time-axis relative to the curves for adapted faces. Fig. 4B illustrates a hypothetical pattern that incorporates such a delay in the development of lateral inhibition: this resembles the data in Fig. 2. Thus, with some important modifications, our results may be consistent with the priming model of Huber and O'Reilly (2003).

How does the outcome of our experiment relate to our first aim, to determine if selective response suppression can account for perceptual-shift face aftereffects? These results present a challenge for such explanations. First, the fact that aftereffects in the responses to both adapted and unadapted faces are found differs from models that postulate that adaptation suppresses responses to adapted stimuli and does not affect responses to unadapted stimuli. Second, most perceptual-shift studies report a *repulsive* aftereffect, i.e., biased perception of test stimuli away from the adapting stimuli, (Butler et al., 2008; Fang & He, 2005; Fox & Barton, 2007; Jiang, Blanz, & O'Toole, 2006; Rhodes & Jeffery, 2006), typically using an adaptation period of 5 s, but also seen with periods ranging from 1–16 s (Leopold et al., 2005; Rhodes et al., 2007). These perceptual shifts are explained as an imbalance in the net population response following suppression of the responses to the adapted face. This imbalance should be reflected in greater elevation of thresholds for the adapted face than for unadapted faces. However, in our first experiment we found the reverse, that thresholds for unadapted faces are even more elevated than thresholds for adapted faces. Therefore, the balance between the adapted and the unadapted faces should have predicted a greater likelihood of perceiving the adapted face in studies using perceptual-shift aftereffects, i.e., an 'attractive' aftereffect, not the 'repulsive' aftereffect documented in studies using perceptual-shift methods.

Before concluding that the current explanation of perceptual-shift aftereffects does not appear to be adequate, one possible and trivial explanation of this discrepancy is that our stimulus set is anomalous and actually does not produce repulsive aftereffects at 5 s of exposure, in contrast to the robust repulsive aftereffects reported by prior studies (e.g., Butler et al., 2008; Fang & He, 2005; Fox & Barton, 2007; Fox et al., 2008; Rhodes & Jeffery, 2006). Therefore, in Experiment 2 we measured perceptual-shift aftereffects with the face stimuli of Experiment 1 to confirm that these generate repulsive face aftereffects similar to those reported by others.

3. Experiment 2: perceptual-bias face aftereffects

Perceptual-bias aftereffects for faces have been reported in the literature for a variety of facial attributes and adapting durations.

Aftereffect magnitudes tend to get larger with increased adapting duration, and whenever any aftereffect is observed the direction of the perceptual bias is always 'repulsive', i.e., away from the adaptor. This is inconsistent with the prediction offered by the data of Experiment 1: these showed that, for a wide range of adapting durations, the suppression of the response was larger for unadapted faces than for the adapted face itself. This predicts 'attractive' perceptual-bias aftereffects. In addition, at around 5 s adaptation, where a number of prior studies observed robust repulsive aftereffects, the difference in response suppression is virtually gone (Fig. 2), which would predict that perceptual-bias aftereffects should be non-existent at this duration.

Following the methods in prior reports (Butler et al., 2008; Fox et al., 2008), we measured the effect of adaptation on the perceived identity of an ambiguous test face. We obtained morph series between pairs of faces used in Experiment 1. The middle range of these morph series contain faces with maximal ambiguity as to their identity, which are used as test stimuli. Observers briefly viewed the ambiguous test stimulus and were then required to indicate which one of the constituent pair of faces the test resembled more. We measured how adapting to either of the constituent faces affected the probability of that face being perceived in the ambiguous test face.

Two possible outcomes are a null effect or an attractive aftereffect, which would indicate an idiosyncrasy in the particular face stimuli or trial design used in Experiment 1 as the source of the disagreement between predicted and actual patterns of perceptual-bias aftereffects. Alternatively, if our results replicate the common finding of repulsive aftereffects, this would suggest that selective suppression of adapted representations does not on its own provide an adequate account of repulsive face aftereffects, and would call for other explanations or additional factors or mechanisms of adaptation.

3.1. Methods

3.1.1. Observers

Four observers with normal or corrected-to-normal vision participated (1 male, ages 25–36). All observers were naïve to the purposes of the experiment. The protocol was approved by the review boards of the University of British Columbia and Vancouver Hospital, and informed consent was obtained in accordance with the principles in the Declaration of Helsinki.

3.1.2. Apparatus and stimuli

Stimuli were displayed on a Samsung SyncMaster 2220WM LCD monitor at 1680×1050 resolution and refresh rate of 60 Hz, using Superlab 4.0 (<http://www.superlab.com>) stimulus presentation software. We used FantaMorph 4 (<http://www.fantamorph.com>) to create morph series between pairs of the original five faces used in Experiment 1, resulting in a total of 10 distinct morph series. Each morph series contained a total of 41 frames that advanced in 2.5% steps from one constituent face to the other. Only the middle 13 frames of each morph series that extended between the 65%/35% and the 35%/65% morph images were used as test stimuli. Adapting stimuli were the five faces of Experiment 1 that served as the constituent faces of the morph series here.

3.1.3. Procedure

The trial design was similar to that used in Experiment 1. Each trial started with an adapting period (5 s), followed by a 50-ms white noise mask, a 150-ms fixation, a 150-ms blank, a 150-ms test face and finally a choice screen that remained until the observer entered their response. Following protocols used in Fox et al. (2008) the 13 test faces in each morph series were shown twice, once after adapting to each constituent face. The observers were

required to indicate which one of the two constituent faces the test face resembled more by selecting it from a choice screen that displayed the alternatives side by side in a spatially random order (left/right). The observers pressed '1' or '2' in the computer keypad to select the face on the left or right, respectively. No feedback was provided. Each observer completed a total of 260 trials composed of 10 morph series, two adapting conditions per morph series, and 13 test images per morph series in a randomized order.

To familiarize the observers to the experimental procedure, a 10-trial warm-up block that used different faces from those used in the experiment was completed prior to starting the experimental session. The responses in the warm-up block were not recorded.

3.1.4. Data analysis

An identity aftereffect score for each morph series was calculated as the difference between the proportion of trials on which the second constituent face was chosen after adapting to the first constituent face, and the proportion after adapting to the second constituent face. Repulsive aftereffects would be reflected in a positive value for the aftereffect score as this would indicate a higher frequency of responding with the second constituent face after adapting to the first, compared to responding after adapting to the second face. In contrast, an attractive aftereffect would be reflected in a negative value. Aftereffect scores were entered into a repeated measures ANOVA to test for any main effects of the morph series. This was followed by a two-tailed *t*-test that checked whether aftereffect scores differed significantly from zero.

4. Results

There was no main effect of morph series, as shown by a repeated measures ANOVA on aftereffect scores ($F(9, 27) = 1.32$, $p > 0.2$). Therefore aftereffect scores were pooled across all subjects and morph series. This yielded an average aftereffect score of 0.15, which was significantly different from zero ($p \ll 0.01$, two-tailed *t*-test). These results indicate a repulsive identity aftereffect in which a 5 s adaptation to one face produced a 15% decrement in the likelihood of seeing that identity in an ambiguous test face compared to adapting to the other face.

In this experiment, we used the stimuli of Experiment 1, and stayed close to that experimental design so as to make sure that our particular face stimuli, or our design where all trial types, e.g., adapting faces, are randomly interleaved, did not in some way eliminate the standard finding of repulsive face aftereffects of other studies. Since this confirmed that our stimulus set does produce repulsive face aftereffects in a perceptual-bias method, the pattern of threshold elevation results of Experiment 1 obtained with the same stimuli indicated that the current explanation of perceptual-shift aftereffects does not appear to be adequate. To confirm this and then to explore alternative explanations, we implemented a simple population-coding model to estimate the perceptual shift predicted from our contrast threshold elevation data.

4.1. Modeling face aftereffects

Our model (Appendix) is based on that developed by Gilbert and Wiesel (1990) to relate tuning changes in orientation-selective cells in cat striate cortex to the perception of tilt illusion. Similar models have been used by others to account for tilt (Clifford et al., 2000) and motion (Kohn & Movshon, 2004) aftereffects, and are based on the activity of multiple units, each selectively responsive to a narrow range of a stimulus attribute, such as motion direction or orientation, and each labeled with their preferred stimulus value. Analogously, we have implemented multiple hypo-

thetical face-selective units, covering a similarity- or identity-axis. For simplicity we represent this in a single-dimension, although the results can be generalized to the multiple dimensions hypothesized for face-space (Valentine, 2001). This model can provide predictions for the pattern of perceptual-shift aftereffects, given our data on threshold changes.

The key aspect of the model is the use of multiple overlapping narrowly-tuned identity-selective units defined on a similarity-based continuous identity-axis (Fig. 5A). In the context of our experiment, one of the units in this structure will be maximally responsive to the adapting face. Unadapted faces are represented by units to either side at a distance dependent on the similarity between the adapting face and the other faces. The 'perceived identity' is based on the population response, which is computed as the response-weighted average of all units, or equivalently, the center of mass of the response distribution.

Let us consider the perceptual shift predicted by our model due to the response suppression found in our data, following adaptation for 1–6.4 s. Our data (Fig. 2) show that in this period, the responses of the adapted unit will be suppressed, but the responses of other units will be suppressed even more. The response of the face-selective units after such adaptation is seen in the left panel of Fig. 5B. The right panel shows the perceptual shift as a function of the distance between adapting and test stimuli after adaptation: deviation towards lower-left and upper-right quadrants denotes repulsive shifts, and deviation towards upper-left and lower-right quadrants denotes attractive shifts. As suggested above, the model predicts an attractive aftereffect, contrary to the repulsive aftereffects documented in the literature. This refutes an account of perceptual face aftereffects based solely on response suppression and indicates that there must be other adaptive processes that counter the effects of response suppression.

Indeed, others have suggested that multiple adaptation mechanisms might exist. One such mechanism is a *sharpening* of stimulus tuning (Desimone, 1996; Grill-Spector, Henson, & Martin, 2006; Wiggs & Martin, 1998). According to the *sharpening* model, adaptation increases the selectivity of the unit that is most responsive to the adapting stimulus; a narrowing of tuning that is due to greater suppression of the responses of this unit to similar stimuli than to the preferred stimulus, resulting in a 'sharper' representation. While response reduction might be the predominant mechanism of adaptation in earlier stages of the visual stream (Albrecht, Farrar, & Hamilton, 1984; Kohn & Movshon, 2004; Movshon & Lennie, 1979), a sharpening mechanism may be more prevalent in extrastriate cortex. For example, Kohn and Movshon (2004) reported that direction-selective cells in area V5 of macaques show a pattern of adaptation different from the response reduction found in V1: adaptation at the preferred direction of these cells reduced their tuning bandwidth rather than suppressing their peak response. Similarly, Krekelberg, van Wezel, and Albright (2006) showed that adaptation resulted in narrower tuning bandwidth for speed-selectivity in macaque V5.

Recent electrophysiological and functional imaging results have uncovered several cortical areas that are specialized for the processing and perception of faces. These areas are mostly upstream from those involving the low-level visual processing, including areas in the ventro-temporal (Kanwisher, McDermott, & Chun, 1997; Kanwisher & Yovel, 2006; McCarthy, Puce, Gore, & Allison, 1997) and frontal cortex (Ishai, Schmidt, & Boesiger, 2005; Tsao, Schweers, Moeller, & Freiwald, 2008), suggesting a higher-level visual processing involving faces compared to simple low-level visual attributes such as spatial frequency and orientation. Given the evidence for a sharpening mechanism of adaptation in higher visual areas, it is plausible that sharpening of representations is also produced by face adaptation. We incorporated this idea in a

revision of our model by reducing the bandwidth of the units in the neighborhood of the adapting value, with the maximum reduction occurring at the unit maximally responsive to the adapting stimulus. The tuning widths of units farther from the adapting stimulus to either side remain unchanged. The response profiles of all units after sharpened tuning around the adapting stimulus is shown in the left panel of Fig. 5C, while the right panel depicts the predicted pattern of aftereffects for all test stimuli: this is consistent with a repulsive aftereffect. If we combine sharpening with the suppression effects indicated by our data on contrast thresholds for adapted and unadapted faces, we find that repulsive aftereffects prevail (Fig. 5D). While the precise shape of the aftereffect curve varies with different values of parameters, the repulsive pattern was robust across a wide range of parameter values.

We also examined whether this revised model combining suppression and sharpening mechanisms predicts the documented logarithmic increase in repulsive face aftereffect sizes (Leopold et al., 2005; Rhodes et al., 2007) as adapting times increased from 1 s to 6.4 s. For each adapting duration, the relative magnitudes of response suppression across the face-selective units were qualitatively based on our contrast-threshold data. For all units, responses decrease with increased adapting time. At short adapting durations, unadapted face units are suppressed more than adapted face units. However, the response for the adapted face decreases faster, and reaches the same level as the unadapted face units at the longest adapting duration of 6.4 s. In Fig. 6A, we plot the aftereffect size for one test face (arbitrarily chosen from a neighborhood that produces a large aftereffect) as a function of adapting duration. For comparison, in Fig. 6B, we include a data plot from Leopold et al. (2005), showing facial identity aftereffects as a function of adapting time. The time course of response suppression based on our contrast-threshold data coupled with a sharpening effect in a simple population-coding model parallels the temporal dynamics of perceptual-shift aftereffects.

5. Discussion

In the present study we introduce a novel face aftereffect based on post-adaptation changes in the contrast sensitivity for recognizing faces. We show that the effects of adaptation are complex, with facilitation of recognition of the adapted face at short adaptation durations but suppression at longer durations, while effects on unadapted faces are minimal at short durations but unexpectedly also show suppression at long durations, even greater than that seen for the adapted face. As suggested by others (Huber & O'Reilly, 2003) such effects on unadapted stimuli may constitute evidence of lateral inhibition from adapted to unadapted representations. Such interactions are not typically considered in conventional models of face aftereffects. Furthermore, the relative balance between the suppression of adapted and unadapted faces at long viewing durations (e.g., 5 s) in our contrast-threshold data appeared problematic for explanations of the repulsive aftereffects revealed by perceptual-bias methods – as seen in Experiment 2 as well as many other studies (Butler et al., 2008; Fang & He, 2005; Fox & Barton, 2007; Fox et al., 2008; Jiang, Blanz, & O'Toole, 2006; Rhodes & Jeffery, 2006; Rhodes et al., 2007) – by response reduction alone.

To confirm this in more detail, we turned to a population-coding model of adaptation. Ideally, models of aftereffects should be able to account for results obtained with different experimental methodologies. We studied whether a common model of adaptation analogous to classical models of lower-level visual adaptation (Coltheart, 1971; Graham, 1989; Mather, 1980), could account for both perceptual-bias aftereffects and the contrast threshold aftereffects obtained at 5 s viewing. To take into consideration our contrast-threshold data the model required modification to

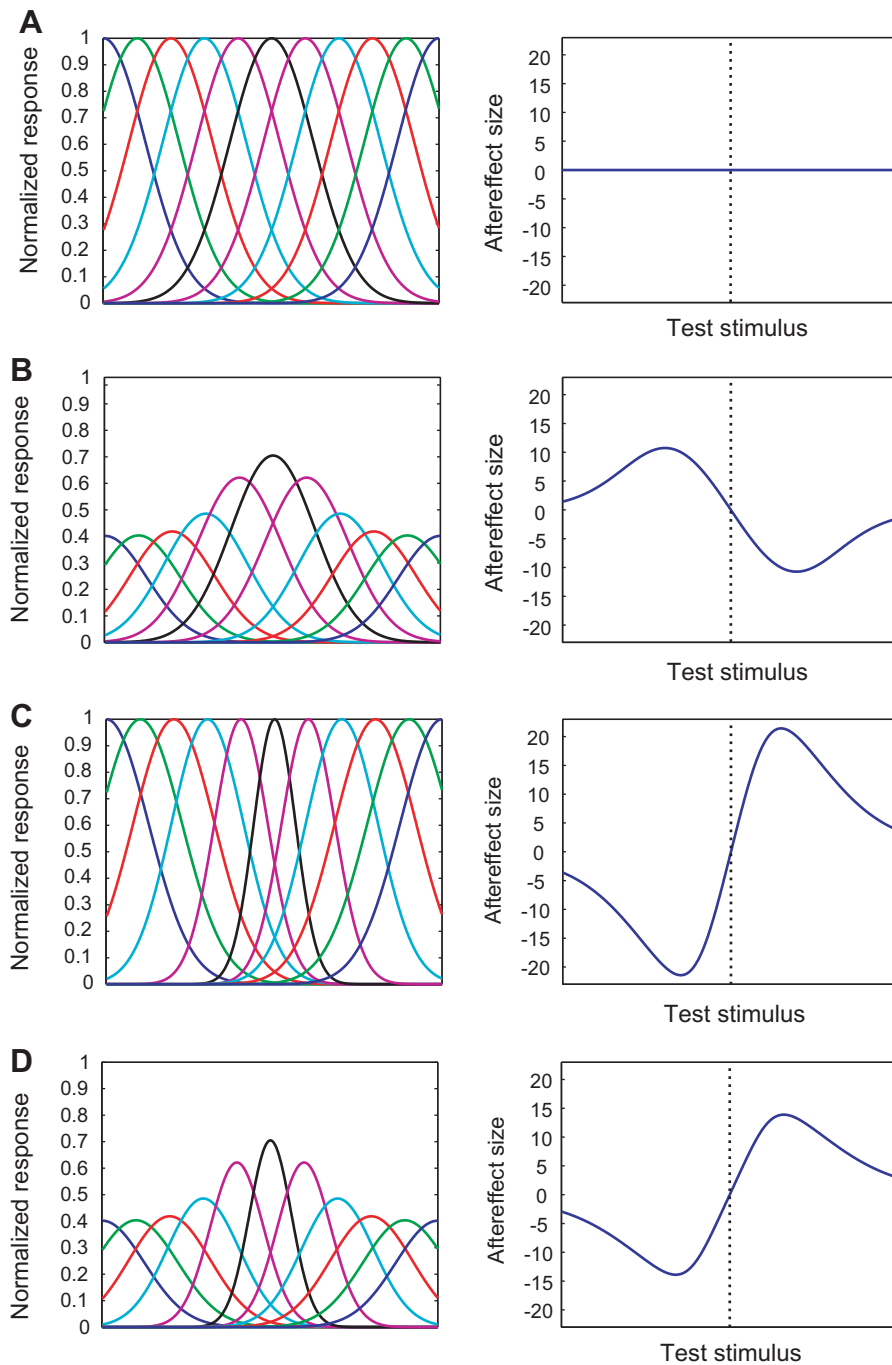


Fig. 5. Model predictions for perceptual-shift aftereffects before adaptation, and after adaptation for three different adaptation schemes. In all graphs, the left panel shows the normalized response profiles of face-selective units with peak tuning values spread along a hypothetical identity axis on the abscissa. The right graph plots the perceptual aftereffect size function of the location of a face stimulus along this identity axis: deformation of the line into the left upper quadrant and right lower quadrant indicate attractive aftereffects, whereas deformation into the left lower quadrant and right upper quadrant indicate repulsive aftereffects. (A) Left panel: response profiles for multiple identity-selective units are plotted at the steady state, prior to any adaptation effects. Right panel: prior to adaptation, the responses of all units are balanced, and the population responses to test stimuli are veridical, producing no aftereffects. (B) Left panel: response magnitudes are reduced as a result of adaptation. In contrast to the conventional models of adaptation where maximum response suppression occurs at the unit tuned to the adapting stimulus (represented by the black line) our model follows the empiric data in Fig. 2 showing that responses away from the adapting stimulus are suppressed more than those at the adapting stimulus value. Right panel: under these conditions, the population response to neighboring test stimuli are shifted towards the adapting value, producing an 'attractive' aftereffect, with the magnitude of the aftereffect dependent on the similarity between adapting and test stimuli. (C) Left panel: In the *sharpening* model of adaptation, the peak responses of the units are unaffected but the bandwidth of the tuning curves are reduced maximally at the adapted location (black line), with the effect lessening for face units situated further from the adapted unit. Right panel: this sharpening effect results in a 'repulsive' pattern of perceptual aftereffects. This aftereffect was consistent across large changes in model parameters: the illustration shows the effects with a 50% decrease in bandwidth at the adapted unit, but similar repulsive aftereffects could still be found with even a 10% decrease. Likewise, repulsive aftereffects were found for a wide range of values for the rate of change of sharpening across units. (D) Left panel: response profiles shown when both response suppression in (B) and the sharpening effects in (C) are combined. Right panel: with combined effects, the net aftereffect profile is a repulsive pattern. This outcome was stable across wide ranges of parameter values, only disappearing when either the magnitude or the window of the sharpening effect were reduced to negligible values.

incorporate suppression of both adapted and unadapted face representations. However, when the pattern of our recognition thresh-

old data was reflected in the model, the result was an 'attractive' perceptual-shift aftereffect, the opposite of the 'repulsive' afteref-

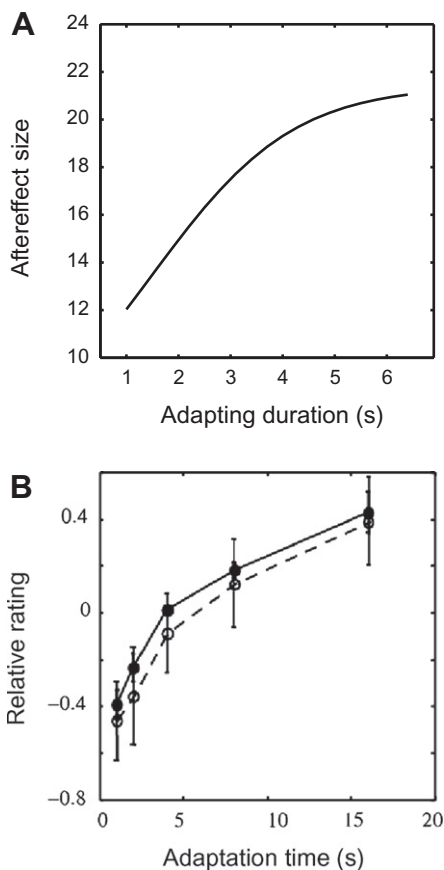


Fig. 6. Temporal dynamics of perceptual face aftereffects predicted by the model. (A) We plot the aftereffect size as a function of adapting duration, as predicted by the model combining response suppression and sharpening effects shown in Fig. 5D, and using the pattern of response suppression for adapted and unadapted units depicted in our data in Fig. 2. At short adapting durations, other faces are suppressed to a greater extent than the same face. This difference diminishes in time, completely disappearing at the end of 6.4 s. The outcome of this dynamic is a repulsive aftereffect that increases approximately logarithmically with adapting duration. (B) Data plot from Leopold et al. (2005) showing the change in face-identity aftereffect size as a function of adapting duration in that study. (Reproduced from Leopold et al., 2005, Proc. Roy. Soc. B.)

fects that have been reported. This suggests that there are additional perceptual effects occurring at 5 s of adaptation. In the model we explored one candidate effect for an additional mechanism, a sharpening of tuning curves of adapted representations, which has been reported for other high-level visual phenomena such as motion perception (Kohn & Movshon, 2004; Krekelberg et al., 2006). Incorporating hypothesized sharpening effects with our suppression data for adapted and unadapted faces in the model generated repulsive perceptual-shift aftereffects and replicated their reported temporal dynamics.

Of course, the findings of the model do not confirm that sharpening is occurring: empiric evidence of its existence is required. Also, the findings do not exclude the possibility that other perceptual effects may explain why repulsive rather than attractive aftereffects occur for face adaptation at durations of 5 s. Rather, the model serves merely to illustrate how selective suppression may not provide a complete picture of adaptation, but may require supplementation by additional factors. Sharpening is simply one possibility, but one that is plausible given its existence for other cortically based adaptation effects, such as motion perception (Kohn & Movshon, 2004; Krekelberg et al., 2006).

The suggestion that adaptation is a complex multi-factorial process is in keeping with the rest of our findings. These show a com-

plex dynamic interplay over a range of durations, with short-term facilitation transitioning to long-term suppression of adapted face representations, possible lateral inhibition of unadapted representations emerging later, and finally a need to incorporate additional mechanisms (sharpening being one possibility) to explain repulsive aftereffects. All of this constitutes a significant elaboration of our view of adaptation as a simple response suppression for the adapted face.

Acknowledgments

This work was supported by CIHR MOP-77615. IO was supported by Natural Sciences and Engineering Research Council, Discovery Grant RGPIN 355879-08. JB was supported by a Canada Research Chair and Michael Smith Foundation for Health Research Senior Scholar Award. This work was presented at the annual meeting of the Vision Sciences Society, Naples 2008.

Appendix A. Model

The response profile, r_i , of face-selective units are modeled as Gaussian functions defined along an identity axis x , as $r_i(x) = \exp(-(x - u_i)^2 / 2\sigma_i^2)$, where the mean u_i and standard deviation σ_i determined the peak tuning and tuning width, respectively for units $i = 1 \dots N$. The responses range between 0 and 1, corresponding to minimum and maximum response. The units were labeled according to their peak tuning. The perceived face identity, F_p , was computed as the response-weighted average of the unit labels as follows:

$$F_p = \frac{\sum_{i=1}^N r_i u_i}{\sum_{i=1}^N r_i}.$$

Response suppression was achieved by scaling the response profile of units by a Gaussian kernel, taking values between peak-response centered on the adapting value, and minimum-response that represented the asymptote of the Gaussian kernel. Peak- and minimum-response values of this kernel could take values between 0 and 1. Sharpening of selectivity was achieved by scaling the standard deviation of the response profile of units by an inverted Gaussian kernel that had an asymptote of one, with the minimum value centered on the adapting value taking on values between 0 and 1. The asymptote of one ensures that the tuning of the units far from the adapted unit were unaffected, while the minimum-value represents the proportion of tuning reduction of the adapted unit. In the absence of empiric data, we chose to make equal the standard deviations of the two kernels determining the width of influence of each effect (response suppression and sharpening of tuning) across the identity-selective units.

The response suppression of all units are modeled as exponential functions of adapting duration where the initial value and the rate of decay was chosen to be higher for the adapted face than the other faces, such that at short adapting durations, the adapted face response was higher than the responses of other faces; however, responses to all faces leveled off to a minimum response at the end of a 6.4 s adaptation period.

References

- Albrecht, D. G., Farrar, S. B., & Hamilton, D. B. (1984). Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *Journal of Physiology*, 347, 713–739.
- Albrecht, D. G., Geisler, W. S., Frazor, R. A., & Crane, A. M. (2002). Visual cortex neurons of monkeys and cats: Temporal dynamics of the contrast response function. *Journal of Neurophysiology*, 88(2), 888–913.
- Andrews, D. P. (1967). Perception of contour orientation in the central fovea. I: Short lines. *Vision Research*, 7(11), 975–997.

- Anstis, S., Verstraten, A. J., & Mather, G. (1998). The motion aftereffect. *Trends in Cognitive Sciences*, 2(3), 111–117.
- Blakemore, C., & Campbell, F. W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, 203(1), 237–260.
- Blakemore, C., & Nachmias, J. (1971). The orientation specificity of two visual after-effects. *Journal of Physiology*, 213(1), 157–174.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436.
- Butler, A., Oruç, I., Fox, C. J., & Barton, J. J. (2008). Factors contributing to the adaptation aftereffects of facial expression. *Brain Research*, 1191, 116–126.
- Carpenter, R. H., & Blakemore, C. (1973). Interactions between orientations in human vision. *Experimental Brain Research*, 18(3), 287–303.
- Clifford, C. W., & Rhodes, G. (Eds.). (2005). *Fitting the mind to the world: Adaptation and after-effects in high-level vision*. Oxford: Oxford University Press.
- Clifford, C. W., Wenderoth, P., & Spehar, B. (2000). A functional angle on some after-effects in cortical vision. *Proceedings of the Royal Society B: Biological Science*, 267(1454), 1705–1710.
- Coltheart, M. (1971). Visual feature-analyzers and after-effects of tilt and curvature. *Psychological Review*, 78(2), 114–121.
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences of the United States of America*, 93(24), 13494–13499.
- Dragoi, V., Sharma, J., Miller, E. K., & Sur, M. (2002). Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nature Neuroscience*, 5(9), 883–891.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman Hall.
- Fang, F., & He, S. (2005). Viewer-centered object representation in the human visual system revealed by viewpoint aftereffects. *Neuron*, 45(5), 793–800.
- Fang, F., Murray, S. O., Kersten, D., & He, S. (2005). Orientation-tuned fMRI adaptation in human visual cortex. *Journal of Neurophysiology*, 94(6), 4188–4195.
- Fox, C. J., & Barton, J. J. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research*, 1127(1), 80–89.
- Fox, C. J., Oruç, I., & Barton, J. J. (2008). It doesn't matter how you feel. The facial identity aftereffect is invariant to changes in facial expression. *Journal of Vision*, 8(3): 11, 1–13.
- Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect, and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology*, 20, 453–467.
- Gilbert, C. D., & Wiesel, T. N. (1990). The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, 30(11), 1689–1701.
- Graham, N. (1989). *Visual pattern analyzers*. New York: Oxford University Press.
- Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, 10(1), 14–23.
- Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science*, 27(3), 403–430.
- Ishai, A., Schmidt, C. F., & Boesiger, P. (2005). Face perception is mediated by a distributed cortical network. *Brain Research Bulletin*, 67(1–2), 87–93.
- Jenkins, R., Beaver, J. D., & Calder, A. J. (2006). I thought you were looking at me: Direction-specific aftereffects in gaze perception. *Psychological Science*, 17(6), 506–513.
- Jiang, F., Blanz, V., & O'Toole, A. J. (2006). Probing the visual representation of faces with adaptation: A view from the other side of the mean. *Psychological Science*, 17(6), 493–500.
- Jiang, F., Blanz, V., & O'Toole, A. J. (2007). The role of familiarity in three-dimensional view-transferability of face identity adaptation. *Vision Research*, 47(4), 525–531.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17(11), 4302–4311.
- Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, 361(1476), 2109–2128.
- Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology*, 97(5), 3155–3164.
- Kohn, A., & Movshon, J. A. (2004). Adaptation changes the direction tuning of macaque MT neurons. *Nature Neuroscience*, 7(7), 764–772.
- Krekelberg, B., van Wezel, R. J., & Albright, T. D. (2006). Adaptation in macaque MT reduces perceived speed and improves speed discrimination. *Journal of Neurophysiology*, 95(1), 255–270.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects [see comment]. *Nature Neuroscience*, 4(1), 89–94.
- Leopold, D. A., Rhodes, G., Muller, K. M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of the Royal Society B: Biological Science*, 272(1566), 897–904.
- Lundqvist, D., & Litton, J. E. (1998). The Averaged Karolinka Directed Emotional Faces – AKDEF.
- Magnussen, S., & Johnsen, T. (1986). Temporal aspects of spatial adaptation. A study of the tilt aftereffect. *Vision Research*, 26(4), 661–672.
- Mather, G. (1980). The movement aftereffect and a distribution-shift model for coding the direction of visual movement. *Perception*, 9(4), 379–392.
- McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, 9(5), 605–610.
- Movshon, J. A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*, 278(5707), 850–852.
- Muller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, 285(5432), 1405–1408.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America A*, 2(2), 147–155.
- Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision Research*, 46(18), 2977–2987.
- Rhodes, G., Jeffery, L., Clifford, C. W., & Leopold, D. A. (2007). The timecourse of higher-level face aftereffects. *Vision Research*, 47(17), 2291–2296.
- Suzuki, S. (2005). High-level pattern coding revealed by brief shape aftereffects. In C. W. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world: adaptation and aftereffects in high-level vision* (Vol. 2, pp. 135–172). Oxford: Oxford University Press.
- Tsao, D. Y., Schweers, N., Moeller, S., & Freiwald, W. A. (2008). Patches of face-selective cortex in the macaque frontal lobe. *Nature Neuroscience*, 11(8), 877–879.
- Valentine, T. (2001). Face-space models of face recognition. In M. J. Wenger & J. T. Townsend (Eds.), *Computational geometric and process perspectives on facial cognition: Contexts and challenges* (pp. 83–113). Mahwah, NJ: Lawrence Erlbaum Associates, Inc..
- Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, 33(2), 113–120.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428(6982), 557–561.
- Wiggs, C. L., & Martin, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion Neurobiology*, 8(2), 227–233.
- Zhao, L., & Chubb, C. (2001). The size-tuning of the face-distortion after-effect. *Vision Research*, 41(23), 2979–2994.