The dynamics of visual adaptation to faces

David A. Leopold[1*], **Gillian Rhodes**[2], **Kai-Markus Müller**[1] and **Linda Jeffery**[2]

[1]*Max Planck Institut für biologische Kybernetik, Spemannstraße 38, 72076 Tübingen, Germany*
[2]*School of Psychology, University of Western Australia, Perth, WA 6907, Australia*

Several recent demonstrations using visual adaptation have revealed high-level aftereffects for complex patterns including faces. While traditional aftereffects involve perceptual distortion of simple attributes such as orientation or colour that are processed early in the visual cortical hierarchy, face adaptation affects perceived identity and expression, which are thought to be products of higher-order processing. And, unlike most simple aftereffects, those involving faces are robust to changes in scale, position and orientation between the adapting and test stimuli. These differences raise the question of how closely related face aftereffects are to traditional ones. Little is known about the build-up and decay of the face aftereffect, and the similarity of these dynamic processes to traditional aftereffects might provide insight into this relationship. We examined the effect of varying the duration of both the adapting and test stimuli on the magnitude of perceived distortions in face identity. We found that, just as with traditional aftereffects, the identity aftereffect grew logarithmically stronger as a function of adaptation time and exponentially weaker as a function of test duration. Even the subtle aspects of these dynamics, such as the power-law relationship between the adapting and test durations, closely resembled that of other aftereffects. These results were obtained with two different sets of face stimuli that differed greatly in their low-level properties. We postulate that the mechanisms governing these shared dynamics may be dissociable from the responses of feature-selective neurons in the early visual cortex.

Keywords: visual perception; face recognition; adaptation

1. INTRODUCTION

The paradigm of visual adaptation has long been an invaluable asset to the psychologist studying perception, since the diverse aftereffects that result can be highly informative about the brain processes underlying how we see (Gibson & Radner 1937; Blakemore & Sutton 1969; Mather *et al*. 1998). In general, adaptation isolates and subsequently distorts perception of particular stimulus attributes, typically biasing perception towards the opposite of the adapting stimulus. This process is often linked directly to diminished responses of feature-selective neurons in the visual cortex (Blakemore & Campbell 1969; Coltheart 1971; Tolhurst & Thompson 1975; Barlow 1990; Bednar & Mukkulainen 2000). Historically, stimuli demonstrating such aftereffects are impoverished, containing little information other than the relevant adapting feature, such as a particular colour or direction of motion.

Many diverse stimuli are effective in adaptation, and the resulting aftereffects, though they impact on different aspects of perception, have much in common. For example, most aftereffects display interocular transfer if the adapting and test eye are different (Gibson 1937; Wade *et al*. 1993), exhibit storage across blank periods (Spigel 1960; Thompson & Movshon 1978), and are restricted to confined spatial zones in the visual field (Gibson 1937; Anstis & Gregory 1965). Also, aftereffects have a finite duration that depends upon adaptor strength and exposure time (Gibson & Radner 1937; Wolfe 1984; Magnussen & Greenlee 1987; Hershenson 1989).

But how does adaptation to a stimulus predispose the brain towards perceiving its 'opposite'? It is alluring to ascribe aftereffects to the aberrant functioning of sensory neurons that have been previously overstimulated. Orientation aftereffects would then derive from an imbalance among orientation-selective neurons in the primary cortical area V1, and motion aftereffects might similarly arise from direction-selective neurons in the motion-selective middle temporal (MT) cortex. But while there is little doubt that such feature-selective neurons contribute to the expression of aftereffects, their precise role is difficult to pinpoint. The fact that very different stimuli cause similar aftereffects poses a challenge for any theory of adaptation tied to a particular functional architecture (van der Zwan & Wenderoth 1995; Clifford 2002; for a review see Leopold & Bondar 2005). Within the domain of orientation, for example, adaptation to luminance-defined bars, known to activate V1 neurons, or to more complicated stimuli unlikely to activate V1 neurons, causes aftereffects with very similar properties (van der Zwan & Wenderoth 1995; Paradiso *et al*. 1989; Joung *et al*. 2000).

Recently, a new family of 'high-level' aftereffects has been described that involves the perception of faces. Specifically, the prolonged exposure to a face can result in the consistent misperception of subsequently presented faces (Webster & MacLin 1999; Leopold *et al*. 2001; Rhodes *et al*. 2003; Watson & Clifford 2003; Webster *et al*. 2004). The perception of faces, unlike that of simpler stimuli, is thought to involve holistic and parallel analysis, and the brain appears to be highly sensitive both to the shape of local features and the spatial relationships between them (for recent reviews see Peterson & Rhodes 2003). It is also thought to proceed along dimensions of semantic and
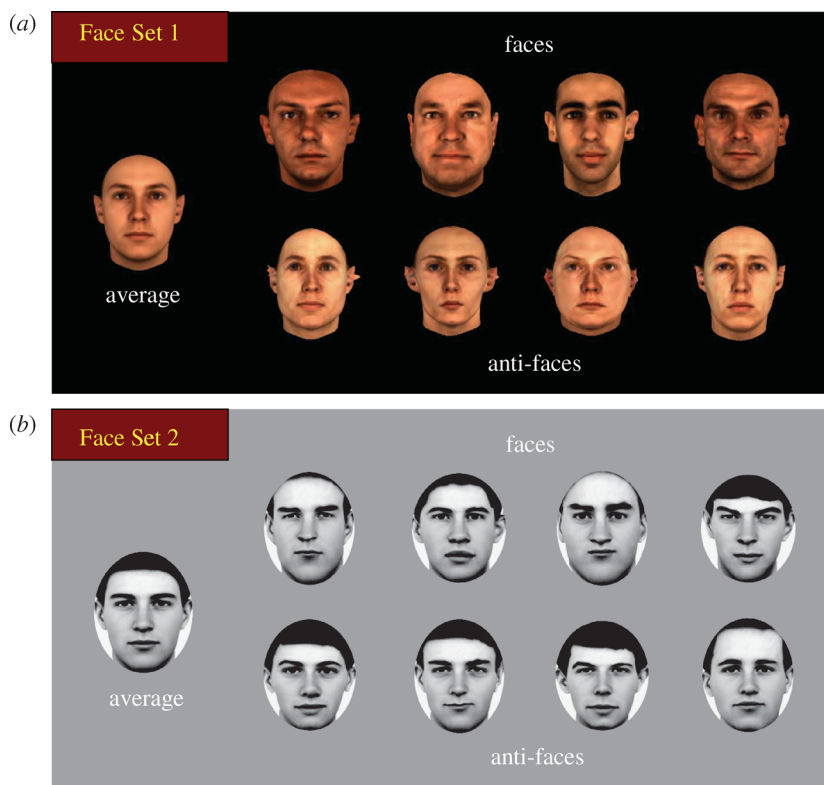
*Author for Correspondence: Unit on Cognitive Neurophysiology and Imaging, National Institutes of Health, MSC 4400, 49 Convent Drive, Building 49, Room B2J-45, Bethesda, MD 20892, USA (leopoldd@mail.nih.gov).

Figure 1. Stimuli used in the present study. (*a*) The MPI faces (Face Set 1), including the average face (left), four authentic faces (top) and their corresponding anti-faces (below). (*b*) The UWA faces (Face Set 2) presented in the same arrangement. A face and its anti-face can be thought of as occupying diametrically opposite locations in a high-level 'face space' with the average face at the centre (Valentine 1991).

social importance, such as identity, expression and attractiveness (Bruce & Young 1998). Several face aftereffects appear to tap directly into these dimensions. In the face identity aftereffect (FIAE), for example, exposure to one face for several seconds systematically distorts the perceived identity of a subsequently viewed different face (Leopold *et al.* 2001; Rhodes *et al.* 2005). The nature of this misperception is not random, but is systematically affected by the specific identity of the adapting face. A neutral test face with average features, for example, is seen as possessing features that are 'opposite' to those of the adapting face, i.e. the adapting face's 'anti-face' is seen (see Leopold *et al.* 2001). This opposite aftereffect appears to conform to established principles of norm-based encoding, where stimuli are encoded not in terms of their absolute structure, but as a deviation from an implicitly stored norm or prototype. Norm-based models have previously been used to account for several aspects of our face perception (Valentine & Bruce 1986; Rhodes *et al.* 1987).

But, given that the attribute of face identity is qualitatively different from those normally identified with simple aftereffects, such as orientation, colour and direction of motion, one might question whether the FIAE is really an aftereffect in the traditional sense. Arguing in favour of this possibility is the fact that the adaptation paradigm for faces is nearly identical to that for established aftereffects, and that both types of aftereffects entail a transient and visible change in the subjective appearance of a subsequent pattern. On the other hand, the aftereffects also differ in some important respects. Unlike with simple stimuli, aftereffects with faces are robust to differences in the size, position and angle of the adapting and test stimuli (Leopold *et al.* 2001; Zhao & Chubb 2001; Rhodes *et al.*

2003; Watson & Clifford 2003). This robustness is more reminiscent of visual object priming, referring to the improved recognition of a visual pattern following previous exposure to that pattern (Bar & Biederman 1998; Biederman & Cooper 1992).

The dynamics of high-level stimuli such as faces remain largely unexplored. With traditional aftereffects, strict relations exist between the magnitudes of the aftereffect, the duration of adaptation and the duration of testing. By contrast, the potency of visual priming is not tightly linked in this way to the prime or test duration. Previous studies using faces have produced aftereffects with both long and short adaptation and test periods (Webster & MacLin 1999; Leopold *et al.* 2001), suggesting that face adaptation occurs over a range of exposure times. But it is unknown whether the relation between these variables resembles that of simple aftereffects. This question is important, as it might provide insight into whether or not the dynamics of these very different aftereffects are dictated by common mechanisms. To examine this issue, we tested the dependence of the FIAE on a wide range of durations of the adapting and test stimuli. To guard against effects that were specific to one particular stimulus set, we performed the same procedure with two different sets of faces and antifaces that were generated by different procedures in different laboratories. The two stimulus sets, while each constructed in the context of a conceptually similar norm-centred *face space* (Valentine 1991; Blanz & Vetter 1999), bore little resemblance to each other, particularly in their low-level properties (see figure 1). In testing the effects of adapting and test duration using the method of constant stimuli, we found that, as with traditional aftereffects, the FIAE grew stronger as a function of adaptation time, and weaker as
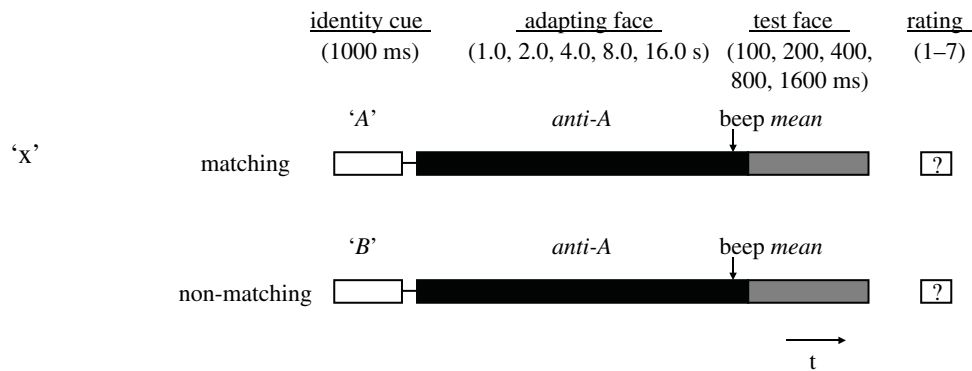
Figure 2. Testing paradigm. Following the presentation of a text cue, the adapting face was shown for a variable time. At 250 ms prior to replacement of the adapting with the test face, a tone sounded briefly to alert the subject to the impending stimulus. Following removal of the test face, the subject was required to rate the degree to which the final percept matched the cued identity. In the UWA data, there were both match and mismatch trials, according to the correspondence between the cue and the anti-face.
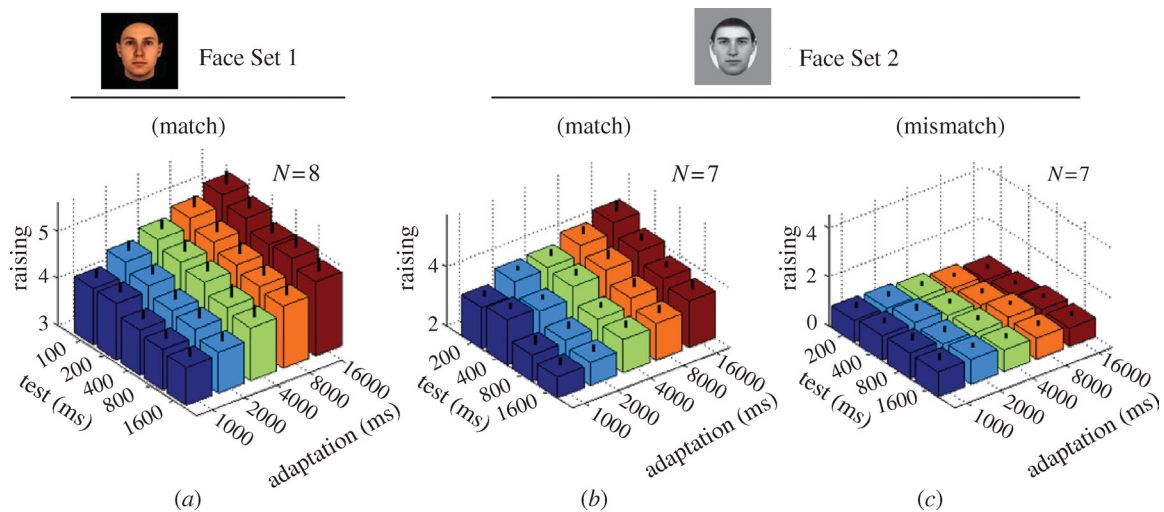

Figure 3. Absolute ratings as a function of test and adaptation times. (*a,b*) In the match conditions, the name cue presented to the subject matched the identity expected following anti-face adaptation. This was done with both face sets. (*c*) In the mismatch condition, the cue name was for a face that did not correspond to the anti-face on that trial. Note the overall lower ratings in the mismatch condition.

a function of test duration. More importantly, we found that the subtle aspects of these dynamics, such as the power-law relationship between the adapting and testing stimulus durations, also resembled that of other after-effects. Based on the similar temporal dynamics of low- and high-level aftereffects, we conclude that the mechanisms underlying their build-up and decay are unlikely to be localized in any particular level of visual processing, and speculate that they may instead reflect the operation of (i) circuits with common temporal processing dynamics in different visual areas, (ii) large-scale networks coordinating representations at different levels of visual analysis, or even (iii) non-sensory mechanisms involved in selective attention and other cognitive processes.

2. METHODS

Half of the experiments were carried out at the Max Planck Institute (MPI) for Biological Cybernetics in Tuebingen, Germany and the other half at University of Western Australia (UWA) in Perth. The aim was to run parallel experiments with stimulus sets differing in their low-level properties. Fifteen subjects participated altogether (eight at the MPI and seven at UWA, five females in each location). All but one were naive to the goals of the study.

(a) *Stimuli*

Figure 1 shows the two sets of stimuli. Subjects were required to learn four different face identities ('Faces'). On each adapting trial they were adapted to a face opposite one of these faces ('Anti-faces'), and were tested with the 'Average' face (see below for details). Face Set 1 (MPI) consisted of full-colour faces derived from 3D head scans, and then morphed using a computer model (Blanz & Vetter 1999). The faces were presented on a 21 inch monitor and subtended 5° horizontal and 7.5° vertical visual angle. The mean face was the average of 100 male and 100 female scans. Identity trajectories in a high-dimensional 'face space' were defined as connecting an individual face to the average face. Anti-faces were created by morphing the average face 40% of the way along the identity trajectory *away* from the original face. The four face/anti-face pairs were the same as those used in a previous study (Leopold *et al.* 2001).

Face Set 2 (UWA) also consisted of the mean face, individual faces and anti-faces. In this case, however, the faces were greyscale photographs morphed using GRYPHON MORPH, and the average was formed from a pool of 20 male faces (for more details see Rhodes *et al.* 2005). The four *anti-faces* were constructed by morphing the structure of average face away from the target face by 80%. Owing to limitations of the

software the textures' 'colours' could not be correspondingly morphed. In light of this, it was decided to map the textures of the greyscale average onto all stimuli used in this set, including the four target faces, so that the anti-face stimuli would not stand out as obviously unusual. The resulting images were sharpened and placed in an oval mask which hid the outer hairline, but not the inner hairline or face outline. They were presented on a 17 inch monitor and subtended a visual angle of 9.5° horizontal and 12.7° vertical.

(b) Procedure

The subjects were first trained to discriminate faces of diminished identity, starting with 100% and then down to 10% (lower-identity faces not shown). This took the form of a four-alternative 'forced choice', in which the subject had to indicate which of the memorized 'Faces' was shown on each trial. The length of this training (from a few minutes to several 1 h sessions) depended on the individual's prior experience with the face set, which was variable. After achieving consistently high performance on these stimuli, the subjects were then trained to rate their perception of identity (rating task; see below). Nonetheless, each day they were retrained briefly on the forced-choice task with low-identity stimuli until their performance was nearly perfect (e.g. greater than 95% correct responses in 80 faces with a 0.15 identity strength presented for 1000 ms). This served to refresh the subjects' familiarity with the faces and to further ensure that their performance did not decline from session to session.

For the rating task, which formed the basis for all the data reported here, the subjects were told that they would be asked to rate their impression of the identity strength (for a cued identity) of faces displayed for varying durations on a seven-point scale ranging from 1 = No Identity to 7 = Full Identity. 'No Identity', in this case, meant that the test face appeared to have no distinguishing features of the cued identity. Because pilot investigations had shown that perception during the test period can be dynamic, subjects were instructed to rate only their impression at the very end of the test face display interval.

The basic trial structure is shown in figure 2. Each trial was initiated when the participant pressed a button. The name cue appeared in the centre of the screen for 1000 ms, followed by the adapting face (an anti-face). A warning beep sounded for 250 ms before the end of the adaptation time, after which the test stimulus (always the mean face) appeared immediately. A trial could be either a match or a mismatch trial, depending on the relationship between the name cue and the anti-face. With Face Set 1 there were only match trials, whereas with Face Set 2, match and mismatch trials were randomly interleaved in equal proportion. Five adaptation times (1.0, 2.0, 4.0, 8.0 and 16.0 s) were crossed with five test durations (100, 200, 400, 800 and 1600 ms) at MPI and four test durations (200, 400, 800 and 1600 ms) at UWA. The stimuli were presented in pseudorandom order, with an equal number of presentations (10) of each combination of face, adapting duration and test duration, carried out over several sessions. One 'block' of these combinations consisted of 100 trials and required approximately 20 min to complete. A trial started when the subject pressed a button on the button box, at which time a cue name appeared on the screen.

3. RESULTS

Figure 3 shows the mean strength of the identity aftereffect on a scale of 1–7 (see §2) as a function of the adapting and

Table 1. Effects of adaptation time and test duration on mean ratings for Face Set 1 and Face Set 2.
(A two-factor 5×5 repeated-measures ANOVA was used to show significant effects of both adaptation and test duration. The two factors did not interact.)

	Face Set 1	Face Set 2
effect of adaptation time	$F_{4,24} = 7.04$, $p < 0.002$	$F_{4,24} = 23.32$, $p < 0.001$
effect of test duration	$F_{4,24} = 55.58$, $p < 0.001$	$F_{3,24} = 16.73$, $p < 0.001$

test stimulus durations. Figure 3a,b shows the responses in the 'match' condition, for the MPI and UWA face sets, respectively, where the name cue corresponds to the anti-face that was shown. Longer periods of adaptation led to higher identity ratings for each test duration, while longer test durations led to lower ratings. These two effects were prominent and significant for both datasets, as revealed by ANOVA analyses (see table 1). In addition, figure 3c shows the mismatch control condition performed with Face Set 2, where the name cue did not correspond to the adapting anti-face. In that case the ratings were on average much lower, with the largest rating smaller than the smallest rating in the matching case, and little effect of adapt or test duration. Pairwise comparisons between the match and mismatch conditions showed that a significant FIAE was still present at the end of the test period for all but the shortest adapting duration, all t's > 4.44, p's < 0.003. Further testing will be needed to determine how long the FIAE lasts.

The basic trends may be seen more clearly in figure 4a,b, where data from both face sets are collapsed across all test and adapt durations, respectively. Here, each subject's overall mean rating was subtracted from each point prior to combining the data, resulting in relative rating values that could be compared across subjects, without regard to their baseline levels. Figure 4a shows monotonically increasing ratings as a function of adaptation time and demonstrates a striking correspondence in the trends elicited by the different stimulus sets. Figure 4b shows that the ratings are highest when the test stimulus is presented only very briefly. To further explore this trend, mainly for purposes of comparison with traditional aftereffects, we re-plotted the data on semi-log coordinates. This approach has been used previously to characterize the build-up and decay of adaptation to simple stimuli (Magnussen & Johnsen 1986; Hershenson 1989). In figure 4c, the relation between the aftereffect strength and logarithm of the adaptation time is nearly linear, and thereby resembles that observed with tilt (Magnussen & Johnsen 1986) and linear motion (Hershenson 1989). This suggests that the logarithmic build-up for simple aftereffects is also present with adaptation to faces. Note that the mismatch trials, computed relative to the average mismatch which was considerably lower than the average match (see figure 3), shows a slight negative slope. Thus, on these trials, increasing adaptation to the 'wrong' anti-face, resulted in a percept that was less and less likely to be seen as the cued identity. Similarly, the straight lines in figure 4d indicate that the aftereffect decay is nearly exponential over time. Note that this exponential decay is characteristic of aftereffects of motion (Sekuler
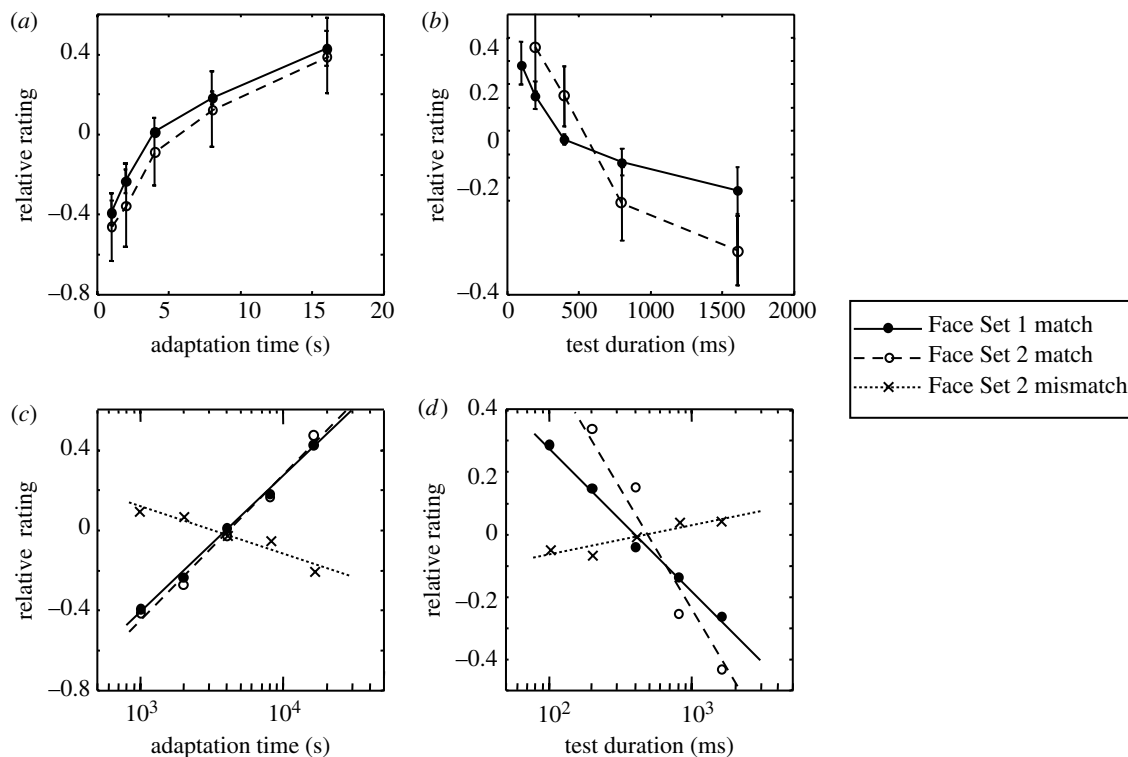
Figure 4. Relative ratings as a function of test and adaptation times for the match condition with Face Set 1 and Face Set 2, and the mismatch condition for Face Set 2. In these plots, the grand mean rating was first subtracted for each subject to emphasize the relative trends as a function of these variables. (*a*) Relative ratings as a function of adaptation time (eight subjects, mean $\pm$ s.e.). (*b*) Relative ratings as a function of test duration (eight subjects, mean $\pm$ s.e.). (*c*) Relative ratings plotted on semi-log coordinates, as a function of adaptation duration. The results are shown for both match and mismatch conditions. See text for details. (*d*) Relative ratings plotted on semi-log coordinates, as a function of test stimulus duration.

1975; Keck & Pentz 1977), orientation (Wolfe 1984; Magnussen & Johnsen 1986; Harris & Calvert 1989) and shape (Krauskopf 1954).

We next considered a more subtle aspect of the adaptation dynamics, again for purposes of comparison to previous work. For traditional aftereffects, the decay *time constant* is not a fixed quantity, as might be expected from some models with simple circuits at their core, but instead varies systematically as a function of the adaptation duration (Taylor 1963; Hershenson 1989, 1993). We thus used variable test durations combined with subjective ratings to explore how the time constant of the FIAE varied as a function of adaptation time.

We plotted subjective ratings against both the adaptation time and test durations, and then extracted iso-rating contours, looking for *combinations* of these two variables that produced a consistent rating. This allowed us to infer the underlying relationship of the variables themselves. The results are shown in figure 5. In figure 5*a*,*b* the iso-rating lines were computed for the two face sets, respectively, based on the same data as in figures 3 and 4. Note that while there are readily identifiable contour lines for each rating level, they are neither straight lines, nor are they similar for different ratings. Instead they appear to have slight curvature and to fan out away from the origin. In other words, the change in test duration for a given increase in adaptation time (i.e. the slope of the contour), depends strongly on which rating contour is being considered.

Yet, when the same data are plotted on log–log co-ordinates (figure 5*c*,*d*), the contours corresponding to

different rating levels are straight and closer to being parallel. The implication of a straight line in log–log coordinates is that the two variables in question obey a power law relationship,

$$y = Ax^b. \tag{3.1}$$

In the present experiments, *x* is the adaptation time, *y* is the test duration, *A* is a constant pertaining to each rating level and *b* is the exponent in the power law. Taking the logarithm of each side, the relationship becomes linear,

$$\log(y) = \log(A) + b \times \log(x). \tag{3.2}$$

The slope *b*, corresponds to the exponent of the power law, and has previously been determined to be approximately 0.5 for other types of aftereffects, corresponding to a square root relationship (Taylor 1963; Lehmkuhle & Fox 1975; Hershenson 1989). In our data, the exponents were higher; approximately 1.5 for the MPI data and 0.8 for the UWA data, and independent of the rating level (i.e. the contours were roughly parallel). The high exponent values indicate that, for adaptation times tested, the increases in the aftereffect duration appeared not to saturate in a way characteristic of simpler aftereffects. In fact, the Set 1 stimuli showed that the aftereffect duration showed the reverse trend, with long adaptation times resulting in disproportionately long-lasting aftereffects. While the implications of this observation are unknown, it might be related to longer-term learning, which has been previously demonstrated following extended periods of adaptation with other stimuli (Wolfe & O'Connell 1986).
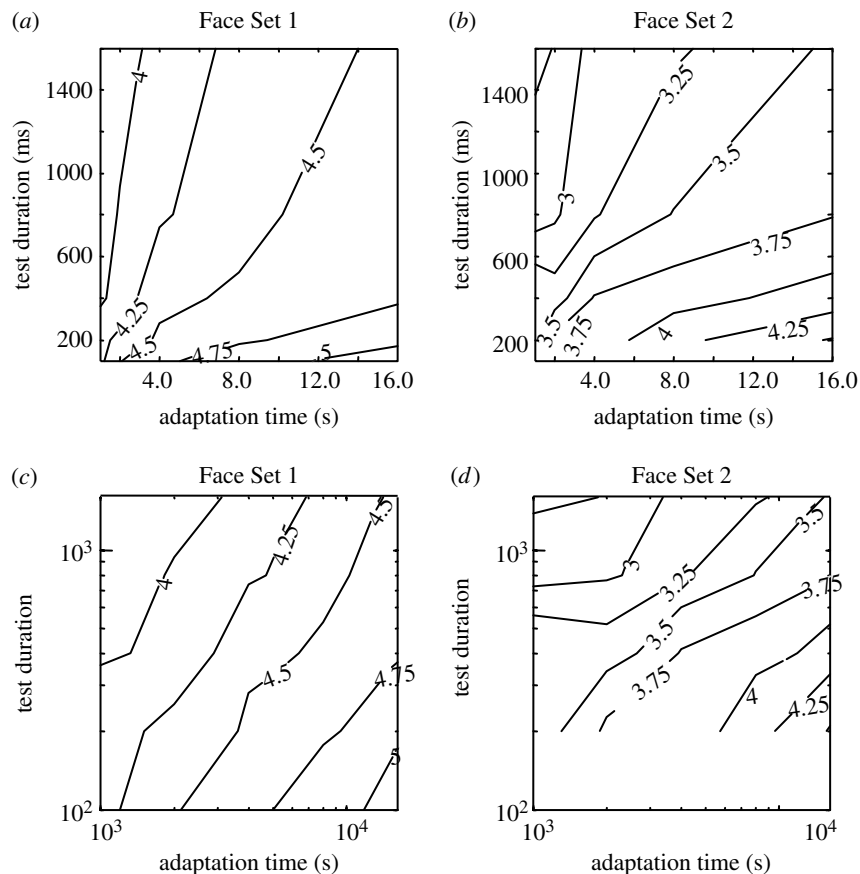
Figure 5. Iso-contour plots showing the combinations of adaptation times and test durations required to achieve a particular grand mean rating. This is shown for both linear and log–log coordinates for both Face Set 1 (*a,c*) and Face Set 2 (*b,d*).

4. DISCUSSION

Taken together, these results show that the FIAE has stereotypic dynamic properties that resemble those described previously for traditional aftereffects. These include logarithmic accumulation with exposure to the adapting stimulus, exponential decay over the test period and a power-law relationship between the adapting and decay durations. This was true for both sets of faces, even though their low-level features were very different, arguing that the observed effects do not critically depend either on the specific morphing algorithm used to generate the average- and anti-faces, or on low-level features of the faces such as colour or texture. These 'classical' dynamics bolster the view that the FIAE is an adaptational aftereffect in the traditional sense, rather than a fundamentally different perceptual effect related to, for example, object priming. The similarities to classical aftereffects are accompanied by some subtle differences, such as the value of the power law relating the two time constants, and the significance of this is unknown.

These results underscore the difficulty in pinpointing the neural locus of aftereffects in general. Given their positional invariance, face aftereffects are unlikely to derive exclusively from circuits in the primary visual cortex. At the same time, their dynamics suggest a high degree of mechanistic overlap with simpler aftereffects, which are often considered to have their origins in early retinotopic processing. The question therefore arises, how is it possible that such diverse aftereffects, apparently resulting from adaptation of distinct populations of neurons in the visual cortex, share temporal properties to such a degree?

One possibility is that the wiring among visual neurons throughout the brain is sufficiently stereotyped that the same dynamics arise wherever competing groups of neurons become differentially adapted. This notion might allow for the present results to fit into a scheme not so different from traditional accounts of simple aftereffects (Coltheart 1971). Such notions normally invoke antagonistic connectivity between a pair of alternative stimulus representations, which could be orientation-selective neurons in V1 or face-selective neurons in the inferotemporal cortex, with the circuit dynamics generated locally in each case. While this may be a parsimonious explanation for the observed results, there is little evidence for it.

A second possibility might be that aftereffects are, by nature, a product of interactions between different processing stages in the brain. With respect to the cortical hierarchy, a perceptual aftereffect might never be accurately described as purely 'low-level' or 'high-level', since very different visual stimuli adapt the same, multiple processing stages, albeit in different ways. Clearly, some aftereffects bear a signature of early processing (e.g. retinotopically restricted adaptation fields) and others of late processing (e.g. invariance to scale and position). But it may be that the other shared aspects of their phenomenology, such as their temporal dynamics, can be attributed to a stereotypic activation of the entire visual cortex that is independent of the specific stimulus. It is

important to consider that regardless of the complexity of a stimulus, at some level the brain delivers a similar product: a subjective impression of visible features of which one can judge colour, brightness, orientation and size. It may be that the general mechanisms underlying this experience, rather than the neurons dedicated to processing its specific features, effectively set the dynamic properties of aftereffects. Regional processing might instead shape the degree of invariance, or even norm-based organization, of aftereffects, without determining the dynamics. While the evidence for this hypothesis is also minimal, recent neuroimaging studies do verify that large-scale networks in the brain, at different cortical processing stages, are affected by periods of prolonged stimulus adaptation (Taylor *et al.* 2000; Tolias *et al.* 2001).

A third, intriguing possibility is that the dynamics of adaptation are not at all determined by the feature-selective neurons that respond to and 'represent' the adapting stimulus, but are instead generated external to the visual system proper. While this possibility is largely unexplored and may be counterintuitive, it does provide an alternative account of the shared pattern of aftereffect build-up and decay demonstrated for very different visual patterns. It might also provide an important link to other phenomena, such as multi-stable perception, for which the specific role of adaptation is long debated, and whose dynamics are thought to be governed by mechanisms related to selective attention (Lumer *et al.* 1998; Leopold & Logothetis 1999). In this vein, it is conceivable that the common aftereffect dynamics similarly reflects intervention from processes related to higher aspects of cognition, such as attention and/or memory.

While it is currently difficult to distinguish among these and other possibilities, the present results strongly suggest that the perceptual distortions following adaptation to faces are aftereffects in the traditional use of the term. Clarifying the neural basis of these common dynamics may provide insight into mechanisms by which the brain, over short time periods, can adjust its sensory processing to meet the requirements of the world. Such processes may be an important element of natural vision.

REFERENCES

Anstis, S. M. & Gregory, R. L. 1965 The aftereffect of seen motion: the role of retinal stimulation and of eye movements. *Q. J. Exp. Psychol.* **17**, 173–174.

Bar, M. & Biederman, I. 1998 Subliminal visual priming. *Psychol. Sci.* **9**, 464–469.

Barlow, H. B. 1990 A theory about the functional role and synaptic mechanisms of visual aftereffects. In *Vision: coding and efficiency* (ed. C. Blakemore), pp, 363–375. Cambridge University Press.

Bednar, J. A. & Mukkulainen, R. 2000 Tilt aftereffects in a self-organizing model of the primary visual cortex. *Neural Comput.* **12**, 1721–1740.

Biederman, I. & Cooper, E. E. 1992 Size invariance in visual object priming. *J. Exp. Psychol.: Hum. Percept. Perform.* **18**, 121–133.

Blanz, V. & Vetter, T. 1999 A morphable model for the synthesis of 3d faces. *Proceedings of SIGGRAPH '99*. New York: ACM Press, pp. 187–194

Blakemore, C. & Campbell, F. W. 1969 On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images. *J. Physiol.* **203**, 237–260.

Blakemore, C. & Sutton, P. 1969 Size adaptation: a new aftereffect. *Science* **166**, 245–247.

Bruce, V. & Young, A. 1998 *In the eye of the beholder: the science of face perception*. New York: Oxford University Press.

Clifford, C. W. 2002 Perceptual adaptation: motion parallels orientation. *Trends Cogn. Sci.* **6**, 136–143.

Coltheart, M. 1971 Visual feature-analyzers and aftereffects of tilt and curvature. *Psychol. Rev.* **78**, 114–121.

Gibson, J. J. 1937 Adaptation, aftereffect, and contrast in the perception of tilted lines. II. Simultaneous contrast and the areal restriction of the aftereffect. *J. Exp. Psychol.* **20**, 553–569.

Gibson, J. J. & Radner, M. 1937 Adaptation, aftereffect, and contrast in the perception of tilted lines. I. Quantitative studies. *J. Exp. Psychol.* **20**, 453–467.

Harris, J. P. & Calvert, J. E. 1989 Contrast, spatial frequency, and test duration effects on the tilt aftereffect: implications for underlying mechanisms. *Vision Res.* **29**, 129–135.

Hershenson, M. 1989 Duration, time constant, and decay of the linear motion aftereffect as a function of inspection duration. *Percept. Psychophys.* **45**, 251–257.

Hershenson, M. 1993 Linear and rotation motion aftereffects as a function of inspection duration. *Vision Res.* **33**, 1913–1919.

Joung, W., van der Zwan, R. & Latimer, C. 2000 Tilt aftereffects generated by bilaterally symmetrical patterns. *Spat. Vis.* **13**, 107–128.

Keck, M. J. & Pentz, B. 1977 Recovery from adaptation to moving gratings. *Perception* **6**, 719–725.

Krauskopf, J. 1954 The magnitude of figural aftereffects as a function of the duration of the test period. *Am. J. Psychol.* **67**, 684–690.

Lehmkuhle, S. W. & Fox, R. 1975 Effect of binocular rivalry suppression on the motion aftereffect. *Vision Res.* **15**, 855–859.

Leopold, D. A. & Bondar, I. 2005 Adaptation to complex visual patterns in humans and monkeys. In *Fitting the mind to the world: adaptation and aftereffects in high-level vision* (ed. C. W. Clifford & G. Rhodes). Oxford University Press, 189–211.

Leopold, D. A. & Logothetis, N. K. 1999 Multistable phenomena: changing views in perception. *Trends Cogn. Sci.* **3**, 254–264.

Leopold, D. A., O'Toole, A. J., Vetter, T. & Blanz, V. 2001 Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neurosci.* **4**, 89–94.

Lumer, E. D., Friston, K. J. & Rees, G. 1998 Neural correlates of perceptual rivalry in the human brain. *Science* **280**, 1930–1934.

Magnussen, S. & Greenlee, M. W. 1987 Saturation of the tilt aftereffect. *Vision Res.* **27**, 1041–1043.

Magnussen, S. & Johnsen, T. 1986 Temporal aspects of spatial adaptation. A study of the tilt aftereffect. *Vision Res.* **26**, 661–672.

Mather, G., Verstraten, F. & Anstis, S. 1998 *The motion aftereffect: a modern perspective*. Oxford University Press.

Paradiso, M. A., Shimojo, S. & Nakayama, K. 1989 Subjective contours, tilt aftereffects, and visual cortical organization. *Vision Res.* **29**, 1205–1213.

Peterson, M. A. & Rhodes, G. 2003 *Perception of faces, objects, and scenes: analytic and holistic processes*. Oxford University Press.

Rhodes, G., Brennan, S. & Carey, S. 1987 Identification and ratings of caricatures: implications for mental representations of faces. *Cogn. Psychol.* **19**, 473–497.

Rhodes, G., Jeffery, L., Watson, T. L., Clifford, C. W. G. & Nakayama, K. 2003 Fitting the mind to the world: face adaptation and attractiveness aftereffects. *Psychol. Sci.* **14**, 558–566.

Rhodes, G., Robbins, R., Jaquet, E., McKone, E., Jeffery, L. & Clifford, C. W. 2005 Adaptation and face perception. In *Fitting the mind to the world: adaptation and after effects in high-level vision* (ed. C. W. Clifford & G. Rhodes). Oxford University Press, 213–240.

Sekuler, R. 1975 Visual motion perception. In *Handbook of perception: vol. V. Seeing* (ed. E. C. Carterette & M. P. Friedman), pp, 387–433. New York: Academic Press.

Spigel, I. M. 1960 The effects of differential post-exposure illumination on the decay of a movement aftereffect. *J. Psychol.* **50**, 209–210.

Taylor, M. M. 1963 Tracking the decay of the aftereffect of seen rotary movement. *Percept. Mot. Skills* **16**, 119–129.

Taylor, J. G., Schmitz, N., Ziemons, K., Grosse-Ruyken, M. L., Gruber, O., Mueller-Gaertner, H. W. & Shah, N. J. 2000 The network of brain areas involved in the motion aftereffect. *Neuroimage* **11**, 257–270.

Thompson, P. G. & Movshon, J. A. 1978 Storage of spatially specific threshold elevation. *Perception* **7**, 65–73.

Tolhurst, D. J. & Thompson, P. G. 1975 Orientation illusions and aftereffects: inhibition between channels. *Vision Res.* **15**, 967–972.

Tolias, A. S., Smirnakis, S. M., Augath, M. A., Trinath, T. & Logothetis, N. K. 2001 Motion processing in the macaque: revisited with functional magnetic resonance imaging. *J. Neurosci.* **21**, 8594–8601.

Valentine, T. 1991 A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Q. J. Exp. Psychol. A, Hum. Exp. Psychol.* **43A**, 161–204.

Valentine, T. & Bruce, V. 1986 The effects of distinctiveness in recognising and classifying faces. *Perception* **15**, 525–535.

van der Zwan, R. & Wenderoth, P. 1995 Mechanisms of purely subjective contour tilt aftereffects. *Vision Res.* **35**, 2547–2557.

Wade, N. J., Swanston, M. T. & de Weert, C. M. 1993 On the interocular transfer of motion aftereffects. *Perception* **22**, 1365–1380.

Watson, T. L. & Clifford, C. W. 2003 Pulling faces: an investigation of the face-distortion aftereffect. *Perception* **32**, 1109–1116.

Webster, M. A. & MacLin, L. H. 1999 Figural aftereffects in the perception of faces. *Psychon. Bull. Rev.* **6**, 647–653.

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

Wolfe, J. M. 1984 Short test flashes produce large tilt aftereffects. *Vision Res.* **24**, 1959–1964.

Wolfe, J. M. & O'Connell, K. M. 1986 Fatigue and structural change: two consequences of visual pattern adaptation. *Invest. Ophthalmol. Vis. Sci.* **27**, 538–543.

Zhao, L. & Chubb, C. 2001 The size-tuning of the face-distortion aftereffect. *Vision Res.* **41**, 2979–2994.