NeuroImage 82 (2013) 426–432

Contents lists available at SciVerse ScienceDirect

NeuroImage

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

Decoding the motion aftereffect in human visual cortex

Hinze Hogendoorn [a,*], Frans A.J. Verstraten [a,b]

[a] Helmholtz Institute, Neuroscience & Cognition Utrecht, Experimental Psychology Division, Utrecht University, Utrecht, The Netherlands
[b] School of Psychology, Faculty of Science, The University of Sydney, NSW 2006, Australia

ARTICLE INFO

Article history:
Accepted 8 June 2013
Available online 15 June 2013

Keywords:
Motion aftereffect
fMRI
Multivariate pattern classification
Visual motion processing

ABSTRACT

In the motion aftereffect (MAE), adapting to a moving stimulus causes a subsequently presented stationary stimulus to appear to move in the opposite direction. Recently, the neural basis of the motion aftereffect has received considerable interest, and a number of brain areas have been implicated in the generation of the illusory motion. Here, we use functional magnetic resonance imaging in combination with multivariate pattern classification to directly compare the neural activity evoked during the observation of both real and illusory motions. We show that the perceived illusory motion is not encoded in the same way as real motion in the same direction. Instead, suppression of the adapted direction of motion results in a shift of the population response of motion sensitive neurons in area MT+, resulting in activation patterns that are in fact more similar to real motion in orthogonal, rather than opposite directions. Although robust motion selectivity was observed in visual areas V1, V2, V3, and V4, this MAE-specific modulation of the population response was only observed in area MT+. Implications for our understanding of the motion aftereffect, and models of motion perception in general, are discussed.

Introduction

The motion aftereffect (MAE) is a much-studied visual illusion in which adaptation to motion causes a subsequently presented stationary pattern to appear to move. The stationary pattern appears to move in the opposite direction of the original, physically moving stimulus. For example, after looking at a stationary point in a waterfall, upon fixating the adjacent rocks, those rocks will appear to drift upwards (Addams, 1834; Anstis et al., 1998). The MAE has been argued to result from an imbalance in oppositely tuned motion detectors (Barlow and Hill, 1963; Wohlgemuth, 1911; but see Sutherland (1961) and below). Adapted neurons respond relatively less strongly to a new stimulus than their oppositely-tuned counterparts. As a result, the balance of activity between the two opponent directions favors the unadapted direction, and an illusory motion percept occurs.

The neuronal basis of the MAE has been the topic of substantial recent study. Direction-selective motion-sensitive neurons are present in multiple visual areas (including V1, V2, V3, and V4; Kamitani and Tong, 2006), and numerous brain areas have been shown to be involved in the MAE (Taylor et al., 2000). Nonetheless, the majority of interest has focused on the possible role of human motion area MT+ in generating the illusory motion percept. Work in non-human primates suggests that cells in this area have appropriate characteristics: for example, Petersen et al. (1985) demonstrated direction-specific adaptation of motion-sensitive cells in owl–monkey MT. Tootell et al. (1995) first presented evidence from human observers implicating MT+ activity in the MAE, and demonstrated that the time-course of activation in MT+ mirrors the time-course of the perceptual illusion. He et al. (1998) showed that both the MT+ activation and the perceptual MAE depend on the stationary test pattern being in the same retinal position as the preceding adaptation. Culham et al. (1999) noted that MT+ activation increased even when adaptation and test phases were separated by a storage period. Subsequent studies continued to implicate MT+ in perception of the MAE (Hautzel et al., 2001; Théoret et al., 2002).

Huk et al. (2001) called into question whether MT+ activity actually underlies the perceptual MAE. They argued that attention to the MAE, rather than the MAE per se, explains elevated MT+ activity. In a series of experiments, they demonstrated that when attention is controlled for, no difference in MT+ activation is observed between conditions in which a MAE is perceived and conditions in which it is not. However, Castelo-Branco et al. (2009) noted that in this study, attention was controlled for by having observers carry out a task on concurrent real motion. They argue that this might interfere with any MAE-related signal in MT+. After replicating Huk et al. (2001) using an attention-to-motion stimulus, they demonstrate that MT+ does in fact show elevated activity during the MAE when attention is allocated to a non-motion feature, such as orientation or color. Accordingly, there seems to be considerable evidence for a MAE-related motion signal in MT+.

A net change in cumulative neural activity is just one way in which a neural population might affect motion perception (e.g. Treue et al.,

* Corresponding author at: Helmholtz Institute, Experimental Psychology Division, Utrecht University, Heidelberglaan 2, 3584 CS Utrecht, The Netherlands. Fax: +31 30 253 4511.
E-mail address: j.h.a.hogendoorn@uu.nl (H. Hogendoorn).

1053-8119/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.neuroimage.2013.06.034

2000). However, most computational models of motion perception propose that it is the relative activity of populations of neurons with different direction tuning that underlies the percept of motion (the distribution shift model; e.g. Mather, 1980; Mather and Harris, 1998; Simoncelli and Heeger, 1996; Sutherland, 1961). In support of this view, adapting to motion does selectively adapt subpopulations of neurons in MT+, even under conditions of directed attention (Huk et al., 2001). This suggests that although an increase in mean MT+ activity does not generate the MAE, neural populations in MT+ might nonetheless be involved in generating the illusion.

An important question then arises: if a shift in the relative activity of neural subpopulations in area MT+ causes the illusory motion percept observed in the MAE, how does that pattern of neural activity compare to the neural activity observed in those same populations when observing real motion? In other words, is illusory motion coded in a similar way to real motion? Three general hypotheses can be formulated:

The first possibility is that the same neural populations are active during real and illusory motions in a particular direction. This might be expected in areas involved in integrating motion signals (such as MT+; e.g. Snowden et al., 1991), since the direction of the resulting motion percepts is similar. In this case, neural activity during illusory motion would be comparable to neural activity during real motion in the same direction (Fig. 1C).

A second possibility is that the MAE is the result of an imbalance in oppositely tuned motion detectors. In this case, MAE motion would be expected to be encoded as a relative suppression of activity corresponding to the adapting direction, together with enhancement of the activity coding for the direction of illusory motion (Fig. 1D).

The third possibility is that the motion percept in the MAE results from a different underlying pattern of neural activity than the motion percept elicited by real motion. This would be expected in early motion processing stages, before any integration occurs. However, it might also be expected in later stages if the direction of motion is encoded as the ensemble activity of motion detectors tuned to all directions, rather than as the imbalance of opponent neurons (Treue et al., 2000). In this case, the MAE could result directly from suppression of the adapted motion direction, in line with findings from neurophysiological recordings in non-human primates (Van Wezel and Britten, 2002). Its neural signature would then be quite different from real motion in the same direction (Fig. 1E). This in turn could explain why the MAE (using a stationary test pattern) is easy to distinguish from real motion, despite both containing motion energy in the same direction (Hiris and Blake, 1992).

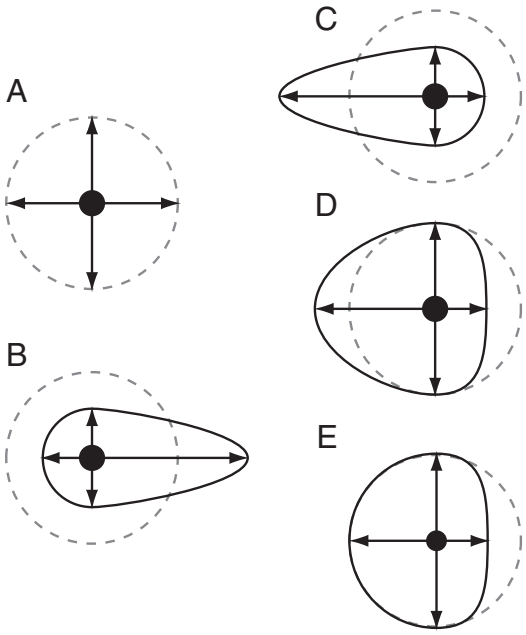

Fig. 1. Schematic illustration of three models of the motion after-effect. Arrows indicate the level of activity of direction-selective neural populations coding for the indicated direction. Ensemble activity is shown with a solid envelope in B–E. A. Response to a dynamic test pattern containing motion in all directions, without previous adaptation to motion. B. Response to an adapting stimulus containing strong rightward motion. C. Hypothetical response during the MAE if the MAE is encoded as the integrated direction of motion. D. Hypothetical response during the MAE if the MAE is encoded as an imbalance of opponent pairs of motion detectors. E. Hypothetical response during the MAE if the MAE is encoded as the ensemble activity of motion detectors tuned to all directions.

The existence of directionally selective subpopulations in human MT+ as well as in earlier visual areas (V1, V2, V3, and V4) can be inferred from functional magnetic resonance imaging (fMRI) studies using pattern classification techniques (Kamitani and Tong, 2006). Using this analysis approach, subtle irregularities in the distribution of direction-selective neurons allow the perceived direction of motion that an observer is viewing to be decoded on the basis of ensemble activity, even when individual voxels contain on the order of a million neurons (Solnick et al., 1984). Importantly, the classification approach works because patterns of neural activity evoked by motion in the same direction are more similar than patterns of activity evoked by motion in different directions. This makes classification performance an effective metric for evaluating the similarity of patterns of neural activity.

Here, we use fMRI in combination with multivariate pattern classification to directly compare the neural activity underlying the illusory motion percept observed during the MAE with neural activity evoked by observing real motion. To do so, we train a classifier on the pattern of fMRI activity in individual visual areas while observers view real motion in one of the four directions. This classifier subsequently categorizes fMRI acquisitions in which observers view a MAE into one of these four directions. The pattern of classification results will allow us to distinguish between these three hypotheses (schematic predictions of each hypothesis are illustrated in Fig. 1).

Methods

Observers

Four observers participated in the experiment. All observers had normal or corrected to normal vision and had previous experience as observers in psychophysical and neuroimaging experiments. All observers gave informed consent before participating. This study was approved by the local ethics committee (Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, Netherlands).

Stimuli

Stimuli were presented on a translucent screen at the head-end of the scanner bore, which the observer viewed at an effective distance of 80 cm by means of a mirror mounted on the head coil. Stimuli were generated and synchronized with the scanner using a PC running Matlab 7.04 (The Mathworks, Inc) with PsychToolbox 2.54 extensions (Brainard, 1997; Pelli, 1997). The scanner room was darkened such that the display was unaffected by ambient light.

Motion stimuli consisted of maximum contrast grayscale random pixel arrays presented in an annulus surrounding a central fixation point, with an inner radius of 2.4° and an outer radius of 10.1°. Both the inner and the outer edge of the annulus were softened with a cosine window of 1.2° wide. Within the annulus, pixel arrays moved in one of the four cardinal directions (up, down, left, and right) at a fixed speed of 3.7°/s. Motion stimuli were presented on a 50% gray background.

Test stimuli consisted of two counter-phasing sparse pixel arrays presented in the same annulus as the motion stimuli (Fig. 2). Each array was presented on a black background, with 20% of pixels within the annulus assigned a random luminance up to 20% of maximum. The two arrays were modulated in sine and cosine phases with a

period of 1 s. The resulting stimulus contained motion energy in all directions at a broad range of speeds. Because the stimulus was presented in counter-phase, it appeared to gently pulsate. Pilot testing showed that this ramping facilitated the perception of a MAE after long storage durations (see Procedure, below). Observers reported perceiving a motion aftereffect for an average of 3.8 s (SD = 0.53 s across observers) during pilot testing outside the scanner, under otherwise identical viewing conditions.

Procedure

Trials were presented in two conditions: a Test condition and a Blank condition. All trials started with 30 s of motion in one of four directions. The motion stimulus was followed by a 12 s storage period, during which the screen was black aside from a central fixation point. This storage period allowed the hemodynamic response to return close to baseline, such that responses to the test stimulus and responses to the adaptor could be separated (He et al., 1998; Huk et al., 2001). Subsequently, in the Test condition the test stimulus was presented for 12 s. In the Blank condition, nothing was presented and the screen remained black for another 12 s. Subsequently, the background was changed to 50% gray for 6 s, followed by the start of the next trial.

Because the test stimulus was presented on a black background, the mean luminance of the two conditions differed during this phase of the trial. This was a necessary consequence of using a dark storage period to maximize the MAE duration (Culham et al., 1999). We anticipated that this would result in a large difference in raw fMRI signal. However, we were interested in patterns of activity reflecting direction-selective subpopulations, and there is no reason to believe that differences in mean luminance would introduce direction-selective effects.

During each trial, observers carried out an attention-demanding change detection task at fixation. Between one and four times in each trial, the central white fixation dot briefly became slightly larger (12.8′ radius instead of 9.1′, for just 17 ms). The task of the observer was to detect and count these changes during the trial, and to use one of the four buttons to report the number when the fixation point turned black at the start of the 6 s delay between trials. Feedback (correct/incorrect) was provided by briefly changing the color of the fixation point to green or red, respectively. The purpose of this task was twofold: to verify that observers were alert and attentive, and also to keep attention engaged at fixation. For each observer, we had verified before the experiment that the task was sufficiently demanding. During the experiment, the four observers correctly responded on 52%, 63%, 78%, and 88% of trials (chance = 25%), indicating that they were paying attention and that they were not performing at ceiling.

Trials were presented in random order in blocks of eight (one trial for each motion direction in each condition). Each observer completed a total of 28 blocks over multiple sessions. Individual sessions took around 2 h, and were distributed over multiple days, spanning up to 6 months.

fMRI acquisition

Data acquisition was performed on a Siemens 3T Trio scanner at the Donders Centre for Cognitive Neuroimaging (Nijmegen, The Netherlands). Functional imaging was conducted using an echo-planar imaging (EPI) sequence (TR = 2000 ms, TE = 30 ms, flip angle = 90°, FOV = 192 mm, matrix size = 64 × 64). 29 oblique slices were acquired for each volume (2.5 × 2.5 mm in-plane resolution, 3.5 mm thick; inter-slice gap = 0.51 mm) covering occipital, temporal, and parietal cortex. Functional runs consisted of 250 functional acquisitions (500 s). The first 4 images of each functional run were discarded to allow for T1 equalization.

High-resolution T1 weighted anatomical scans were also acquired for each subject. All functional runs were co-aligned, motion corrected, and co-registered to anatomical scans. Individual voxel time-courses were low-pass filtered at 0.005 Hz to remove linear trends and scanner drift. No temporal or spatial smoothing was applied.

ROI definition

Localization of MT+ was achieved in each observer with functional data from at least 2 separate localizer runs. In these runs, observers viewed low contrast concentric sinusoidal gratings presented in the same annulus as the motion stimulus. The grating alternately drifted inwards and outwards for 4 s for a total of 16 s, and then remained stationary for 16 s. MT was identified as the contiguous region near the occipital continuation of the inferior temporal sulcus (Dumoulin et al., 2000) that responded more strongly to the moving grating than to the stationary grating.

Retinotopic visual areas V1, V2, V3 (dorsal and ventral), and V4 were similarly identified in at least 2 separate runs, using standard polar angle retinotopic mapping procedures (Sereno et al., 1995) on cortical surfaces flattened using the Mr. Gray toolbox (Wandell et al., 2000).

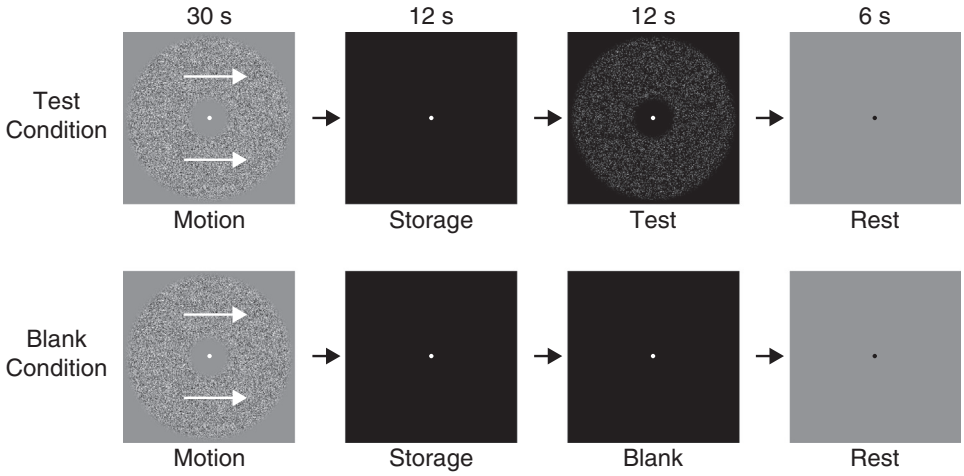

Fig. 2. Experimental timeline. Each trial started with 30 s of motion in one of the four directions, followed by a 12 s storage period during which the screen was dark. Subsequently, a 12 s dynamic test pattern was displayed in the Test condition. In the Blank condition the screen remained dark for 12 s. Trials in both conditions ended with a 6 s gray screen.

fMRI analysis

For each observer, fMRI data for each block were analyzed by training a linear discriminant classifier on data from the remaining 27 blocks (Carlson et al., 2003; Duda and Hart, 1973; Kamitani and Tong, 2006). The classifier was trained to discriminate between trials with different motion direction during the motion phase. Chance performance was therefore 25%. In line with Kamitani and Tong (2006), we considered human motion area MT+ and four retinotopic early visual areas (V1, V2, V3 and V4) using identical analysis procedures.

For motion classification, the classifier was trained on the 15 acquisitions during the motion phase of each trial, delayed by 4 s to compensate for the hemodynamic response function. The classifier was trained separately for each block, using trials from both the Test and Blank conditions from all other blocks. For each block, the classifier was therefore trained on a total of $15 \times 2 \times 27 = 810$ acquisitions per motion direction. Subsequently, we evaluated classification performance at each single time-point during the course of the trial. This allowed us to investigate whether and how activation patterns of real motion generalized to activity patterns associated with illusory motion.

We were particularly interested in the first seconds of the Test phase, since this was the period during which a MAE was reported. The mean MAE duration of 3.8 s corresponded to the first 2 acquisitions (4 s). After compensating for the hemodynamic delay, this corresponded to the period from 46–50 s after the start of the trial. In order to eliminate the possibility that classification performance is based on the trailing hemodynamic response, classification performance is compared in Test and Blank conditions at the same point in time. Any directional response evident in the Test condition, but not the Blank condition, is therefore not attributable to a remnant of the adapting motion.

Results

Visual areas V1–V4

As expected, mean activation was identical on Test and Blank trials during the Motion and Storage phases in each visual area (Fig. 3, top panels). Activation only diverged during the last phase (42–54 s), during which a test pattern was presented in the Test condition but not in the Blank condition.

A classifier trained on the pattern of activation during real motion was able to perform above-chance discrimination of motion direction while observers viewed real motion (0–30 s), in each visual area, replicating Kamitani and Tong (2006). One-sample t-tests over classification performance on all single acquisitions during real motion verified that this effect was highly significant within individual observers ($p < 0.0001$ for all visual areas and all observers). Furthermore, performance was remarkably similar in each visual area, despite considerable differences in the mean number of voxels included in each ROI (368, 245, 186, and 85, for V1–V4, respectively).

When the motion stimulus was removed from the screen, classification assignments essentially reversed. Acquisitions were assigned to the original motion direction with lower-than-chance frequency. Instead, they were assigned to orthogonal or opposite directions. This occurred irrespective of whether a test pattern was presented (and a MAE observed) or not. Although presentation of the test pattern elicited a strong increase in mean activation, it had no impact on classification performance: across visual areas, classification in Test and Blank conditions was indistinguishable throughout the entire trial.

MT+

The classifier trained on the pattern of activation in MT+ during real motion was able to discriminate the direction of motion perceived by the observer, again replicating Kamitani and Tong (2006). One-sample t-tests over classification performance on all single acquisitions during real motion confirmed that this effect was significant for individual observers (all $p < 0.0001$).

It is interesting to note, however, that the pattern of classification errors is different in MT+ than in areas V1–V4. In V1–V4 the classifier was able to distinguish the original direction from all other directions. On the basis of MT+ activation, the classifier performed well at discriminating orthogonal directions of motion, but frequently misclassified motion as the opposite direction (Fig. 4). One-sample t-tests over the proportion of single acquisitions during real motion classified as the opposite direction showed that this effect was significantly above chance within all individual observers (all $p < 0.02$).

Observers perceived a MAE during the first 3.8 s of the test pattern, which was presented in the Test condition, but not in the Blank condition, 42 s into the trial. Accordingly, we were especially interested in any difference in classification between Test and Blank conditions between 46–50 s.

Classification performance in Test and Blank conditions diverged when the test pattern was presented and observers perceived a MAE (Fig. 5). Importantly, the classifier did not classify the pattern of neural activation concurrent with the illusory motion percept as being similar to real motion in the direction opposite to what the observer had viewed on that trial. In other words, illusory motion was not neuronally similar to real motion in the same direction. Instead, the pattern of activity associated with the MAE was more similar to real motion in an orthogonal direction.

A 2×3 factorial repeated measures analysis of variance on the classification assignments 46–50 s after motion onset confirmed a significant interaction effect of condition and motion direction ($F(2) = 8.927$, $p = 0.016$; Fig. 6). Post-hoc paired-samples t-tests indicated that trials in the Test condition (in which a MAE was observed) were significantly less frequently assigned to the direction of real motion observed on that trial ($t(3) = -4.34$, $p = 0.023$), instead being significantly more frequently assigned to an orthogonal direction ($t(3) = 4.30$, $p = 0.023$).

Discussion

We investigated the neural basis of the motion aftereffect using fMRI in combination with multivariate pattern classification. In this way, we directly compared the pattern of neural activity evoked by viewing real motion in one of four directions with the neural activity that is evident when experiencing the resulting motion aftereffect.

We were able to decode the direction of motion which an observer was viewing on the basis of ensemble activity in each of V1, V2, V3, V4, and MT+ individually. After the motion stimulus disappeared, classification assignments reversed, reflecting the adapted state of the neural population. The subsequent presentation of a test pattern induced a MAE. The MAE was associated with a shift in the pattern of neural activity in MT+ (but not in V1–V4). However, this shift was not towards the direction of the perceptual MAE (i.e. opposite to the direction of adapting motion). Instead, it was away from the adapted direction, and towards the orthogonal directions.

The fact that we observe reliable classification performance in areas V1–V4, but observe a MAE effect on classification performance only in MT+, provides convergent evidence that MT+ is uniquely involved in generating the illusory motion that characterizes the MAE. However, the pattern of classification performance during the MAE did not correspond to any of our three predictions. In fact, because both the adapted and the MAE directions were underrepresented in the classification results during the MAE (with orthogonal directions as a result being symmetrically overrepresented), the pattern of fMRI activity in MT+ during the MAE did not uniquely determine the perceived direction of illusory motion (Fig. 6C).

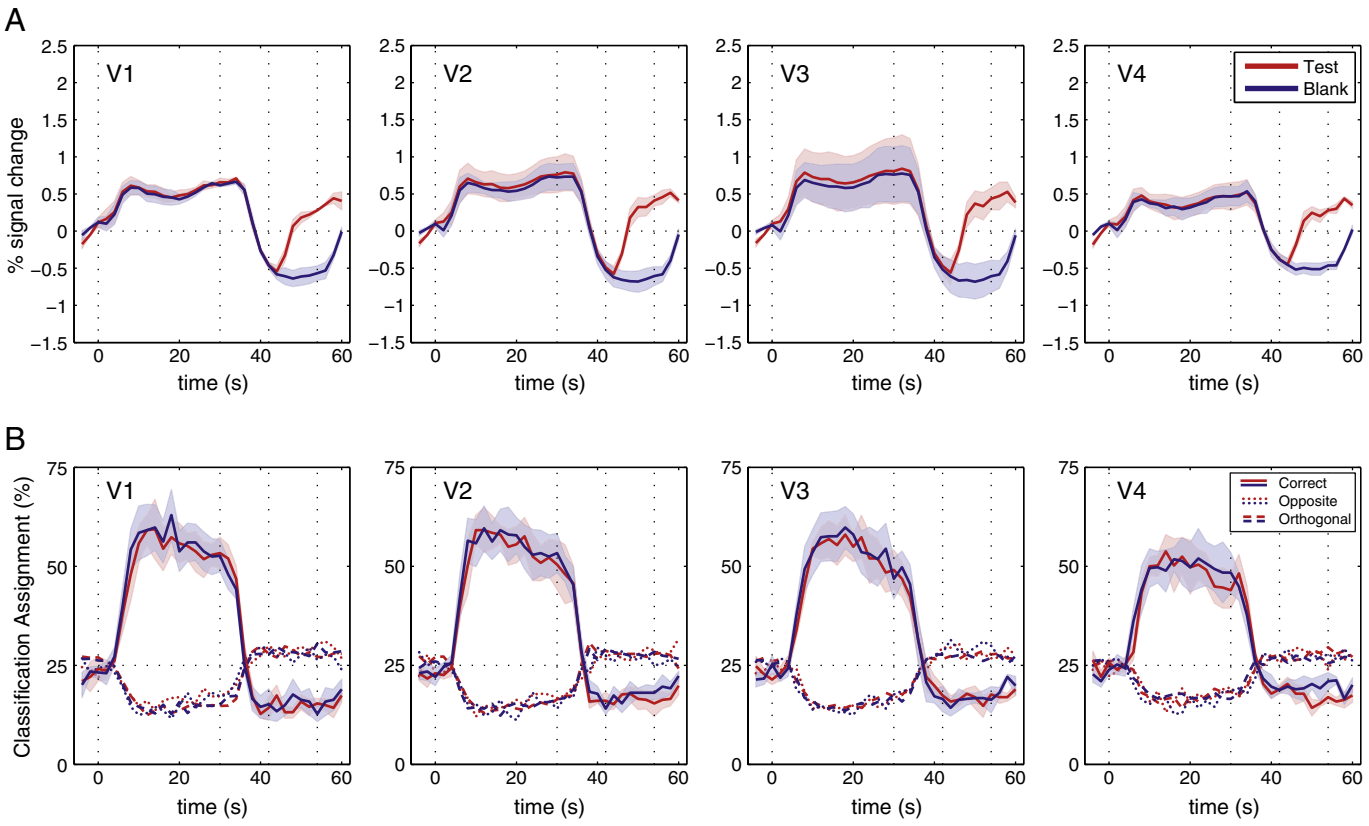

Fig. 3. Motion classification in visual areas V1, V2, V3 and V4. Mean percent signal change (A) and classification assignments (B) as a function of time over the course of the trial. Each trial encompassed 30 s of real motion, followed by a blank period for 12 s and a test pattern for 12 s (in the Test condition only; this period was blank in the Blank condition). Vertical dotted lines indicate the borders of these periods. A. Mean percent signal change. Shaded areas indicate standard errors of the mean across observers. The presentation of a test pattern evokes considerable activity in all four areas. Signal change is otherwise identical for Test and Blank trials. B. Classification performance of a classifier trained on the real motion. Chance performance is 25%. Solid, dotted, and dashed lines denote time points assigned to the true direction of motion, the opposite direction of motion, and an orthogonal direction of motion, respectively. Shaded areas indicate standard errors of the mean across observers; for clarity, these have been omitted from opposite and orthogonal classification. In both Test and Blank conditions, the classifier was able to decode the direction of motion the observer was seeing while motion was presented. When the stimulus was removed from the screen, classification performance dropped below chance. Additionally, although mean activation increases strongly when the test pattern is presented (top panels), this has no consequences on classification (bottom panels).

One possible reason for this finding could be the organization of direction-selective neurons within voxels in MT+. Overall, the pattern of direction classification was remarkably similar in areas V1, V2, V3, and V4. Conversely, although the classifier trained on MT+ activity was also able to decode the perceived direction of motion, the pattern of errors was different from V1–V4. In MT+, but not in V1–V4, real motion

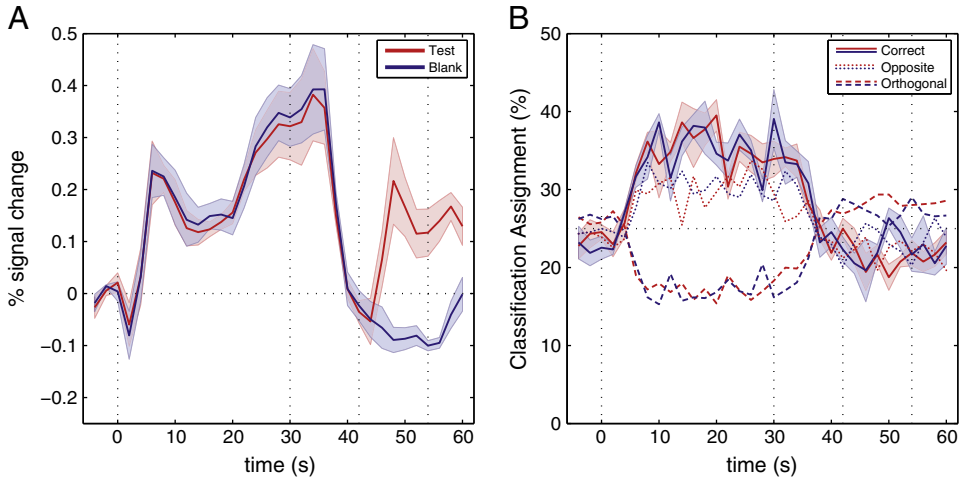

Fig. 4. Motion classification in visual area MT+. A. Mean percent signal change as a function of time over the course of the trial. Shaded areas indicate standard errors of the means across observers. Activity profiles of Test and Blank trials are identical until presentation of the test stimulus, as expected. B. Classification performance as a function of time. Solid, dotted, and dashed lines indicate time points assigned to the true direction of motion, the opposite direction of motion, and an orthogonal direction of motion, respectively. Although the classifier performed above chance at classifying real motion, it frequently misclassified motion in the opposite direction. For individual comparisons of Test and Blank conditions, see Fig. 5.

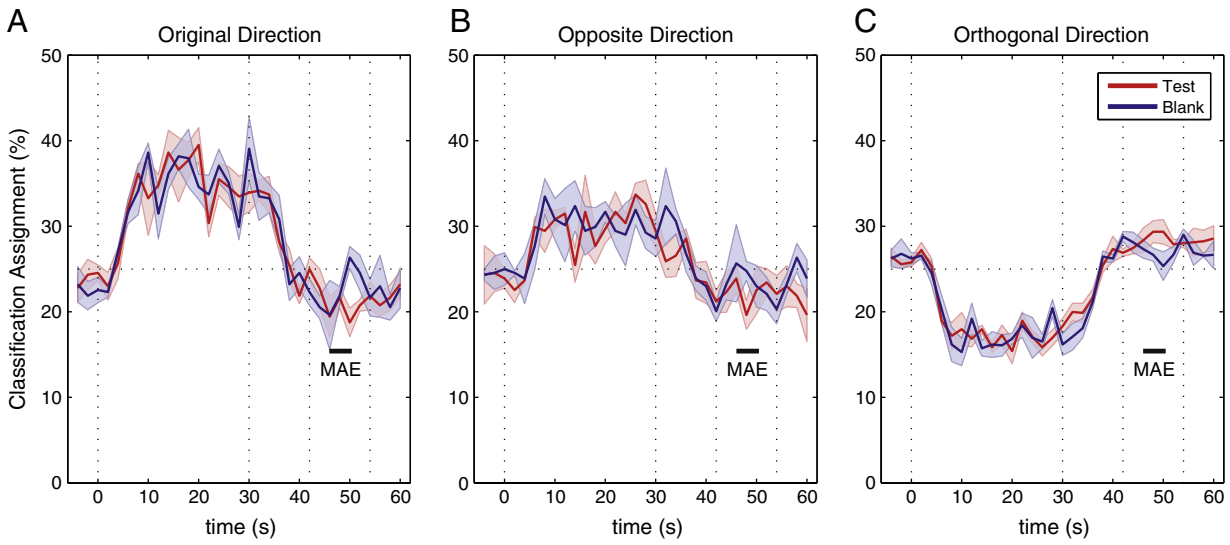

Fig. 5. Mean direction classification as a function of time for a classifier trained on MT+ activity during real motion. Plots show the proportion of trials classified as the original direction (A), the opposite direction (B), or an orthogonal direction (C). Shaded areas indicate standard errors of the mean across observers.

was mis-classified as the direction opposite to the real direction with above-chance frequency (Fig. 5B). This tendency of the classifier to confuse opponent motion directions in MT+ could be the result of direction-selective neurons with opposite direction tuning tending to be close together, and thereby falling in the same fMRI voxel.

Neurophysiological studies in non-human primates have demonstrated a columnar organization of direction-selective neurons in MT+ (e.g. Allbright et al., 1984; Geesaman et al., 1997; Malonek et al., 1994). Importantly, Allbright et al. (1984) observed that sequential macaque MT+ neurons sampled on a single penetration varied continuously in their direction preference, but occasionally jumped by 180°, particularly in tangential penetrations. Optical imaging in owl monkeys similarly revealed that the activation patterns in MT+ produced by motion in opposite directions show significant overlap (Malonek et al., 1994). Geesaman et al. (1997) came to a similar conclusion using 2-deoxyglucose labeling in macaque. Altogether, these studies demonstrate that neurons with opposite direction tuning tend to be close together in monkey MT+. The pattern of classification results based on our fMRI data is therefore consistent with the interpretation that direction-selective neurons might be similarly organized in human MT+.

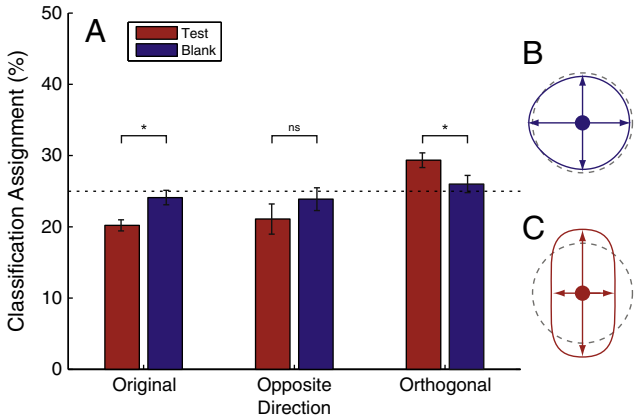

Fig. 6. A: Mean direction classification during the motion aftereffect for a classifier trained on MT+ activity during real motion. Error bars indicate standard errors of the mean. During the Blank condition, essentially no motion information can be decoded: the classifier assigns trials to original, opposite, and orthogonal directions at chance level. However, in the Test condition, systematic motion information can be decoded when a test pattern is presented and a motion aftereffect is perceived: trials are more frequently assigned to the orthogonal motion directions, and less frequently to the original and opposite motion directions. B and C: Classification results plotted in the same way as the model predictions in Fig. 1, for Blank and Test conditions respectively. Results are aligned as if the adapting direction was rightwards.

Another possible reason why the pattern of classification performance differed in MT+ could be its putative motion integration role (e.g. Snowden et al., 1991). When viewing a moving pattern composed of separate components moving in different directions, V1 neurons respond to the direction of individual components, whereas MT+ cells are more frequently selective for the direction of the pattern as a whole (Movshon et al., 1985; Rodman and Allbright, 1989). Determining the direction of motion on the basis of the activity of component neurons requires substantial inhibitory interactions (Snowden et al., 1991). At the voxel level, these inhibitory interactions would amplify differences between interacting components. However, components moving in opposite directions are never part of the same moving surface, and might even activate similar inhibitory interactions to intermediate orthogonal directions. At the voxel level, this could result in the neural activity evoked by opposite directions of motion appearing similar.

Given the classifier's tendency to confuse the patterns of MT+ activation associated with opponent motion directions, our results are most consistent with a distribution-shift model of the MAE (Fig. 1E). This interpretation is also in line with physiology studies in non-human primates showing that adapting to motion suppresses the response of MT neurons tuned to the adapting direction, but does not enhance responses to the opposite direction (Van Wezel and Britten, 2002).

Our results therefore indicate a unique role for MT + in generating the MAE, but the fact that fMRI activation patterns overlap in this area for motion in opposing directions limits the conclusions that can be drawn regarding the neural substrate of the MAE. Nevertheless, taking this into account, our findings are broadly consistent with the interpretation that the illusory motion percept observed during the MAE is driven by a shift in the overall population response in area MT+, rather than by biases in paired, oppositely tuned motion detectors (Grossberg, 1976; Grunewald and Lankheet, 1996; Raymond and Braddick, 1996).

Finally, observed adaptation effects were strikingly robust in all visual areas we investigated: when the moving stimulus was removed from the screen, the performance of a classifier trained on real motion reversed. This reversal of classification performance was maintained for an unexpectedly long time. In fact, in V1–V4 especially, classification performance had not reduced to chance even 30 s after the offset of the motion stimulus. This is not attributable to the hemodynamic delay,

since the mean activation returns to baseline much sooner. Even the presentation of a test pattern, containing real motion in all directions, did not eliminate or reduce the classification bias. Because in the present experiment the next trial started before adaptation had returned to baseline, further research will be necessary to be able to follow the full time-course of recovery from adaptation, which could well be substantially longer than the period which was investigated here (e.g. Lorenceau, 1987; Rose and Lowe, 1982).

Acknowledgments

This study was supported by the Netherlands Organization for Scientific Research (NWO-Pionier to FV). We thank Serge Dumoulin and two anonymous reviewers for their helpful comments.

Conflict of interest

The authors declare no conflict of interest.

References

Addams, R., 1834. An account of a peculiar optical phenomenon seen after having looked at a moving body. Lond. Edinb. Philos. Mag. J. Sci. 5, 373–374.
Allbright, T.D., Desimone, R., Gross, C.G., 1984. Columnar organization of directionally selective cells in visual area MT of the macaque. J. Neurophysiol. 51 (1), 16–31.
Anstis, S.M., Verstraten, F.A.J., Mather, G., 1998. The motion aftereffect: a review. Trends Cogn. Sci. 2, 111–117.
Barlow, H.B., Hill, R.M., 1963. Evidence for a physiological explanation of the waterfall illusion. Nature 200, 1345–1347.
Brainard, D.H., 1997. The psychophysics toolbox. Spat. Vis. 10, 433–436.
Carlson, T.A., Schrater, P., He, S., 2003. Patterns of activity in the categorical representations of objects. J. Cogn. Neurosci. 15, 704–717.
Castelo-Branco, M., Kozak, L.R., Formisano, E., Teixeira, J., Xavier, J., Goebel, R., 2009. J. Neurophysiol. 102, 3016–3025.
Culham, J.C., Dukelow, S.P., Vilis, T., Hassard, F.A., Gati, J.S., Menon, R.S., Goodale, M.A., 1999. Recovery of fMRI activation in motion area MT following storage of the motion aftereffect. J. Neurophysiol. 81, 388–393.
Duda, R.O., Hart, P.E., 1973. Pattern Classification and Scene Analysis. Wiley, New York.
Dumoulin, S.O., Bittar, R.G., Kabani, N.J., Baker Jr., C.L., Le Goualher, G., Pike, G.B., Evans, A.C., 2000. A new anatomical landmark for reliable identification of human area V5/MT: a quantitative analysis of sulcal patterning. Cereb. Cortex 10, 454–463.
Geesaman, B.J., Born, R.T., Andersen, R.A., Tootell, R.B.H., 1997. Maps of complex motion selectivity in the superior temporal cortex of the alert macaque monkey: a double-label 2-deoxyglucose study. Cereb. Cortex 7, 749–757.
Grossberg, S., 1976. Adaptive pattern classification and universal recoding: part I: parallel development and coding of neural feature detectors. Biol. Cybern. 23, 121–134.
Grunewald, A., Lankheet, M.J.M., 1996. Orthogonal motion after-effect illusion predicted by a model of cortical motion processing. Nature 384 (28), 358–360.
Hautzel, H., Taylor, J.G., Krause, B.J., Schmitz, N., Tellmann, L., Ziemons, K., Shah, N.J., Herzog, H., Muller-Gartner, H.W., 2001. The motion aftereffect: more than area V5/MT? : evidence from ^{15}O-butanol PET studies. Brain Res. 892, 281–292.
He, S., Cohen, E.R., Hu, X., 1998. Close correlation between activity in brain area MT/V5 and the perception of a visual motion aftereffect. Curr. Biol. 8, 1215–1218.
Hiris, E., Blake, R., 1992. Another perspective on the visual motion aftereffect. Proc. Natl. Acad. Sci. U. S. A. 89, 9025–9028.
Huk, A.C., Ress, D., Heeger, D.J., 2001. Neuronal basis of the motion aftereffect reconsidered. Neuron 32, 161–172.
Kamitani, Y., Tong, F., 2006. Decoding seen and attended motion directions from activity in the human visual cortex. Curr. Biol. 16 (11), 1096–1102.
Lorenceau, J., 1987. Recovery from contrast adaptation: effects of spatial and temporal frequency. Vis. Res. 27 (12), 2185–2191.
Malonek, D., Tootell, R.B.H., Grinvald, A., 1994. Optical imaging reveals the functional architecture of neurons processing shape and motion in owl monkey area MT. Proc. R. Soc. B Biol. Sci. 258, 109–119.
Mather, G., 1980. The movement aftereffect and a distribution-shift model for coding the direction of a visual movement. Perception 9, 279–392.
Mather, G., Harris, J., 1998. Theoretical models of the motion aftereffect. In: Mather, G., Verstraten, F.A.J., Anstis, S.M. (Eds.), The Motion Aftereffect: A Modern Perspective. MIT Press. Cambridge, MA, pp. 157–186.
Movshon, J.A., Thompson, I.D., Gizzi, M.S., Newsome, W.T., 1985. The analysis of moving visual patterns. In: Chagas, C., Gattass, R., Gross, C. (Eds.), Pattern Recognition Mechanisms. Vatican Press, Rome.
Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat. Vis. 10, 437–442.
Petersen, S.E., Baker, J.F., Allman, J.M., 1985. Direction-specific adaptation in area MT of the owl monkey. Brain Res. 346 (1), 146–150.
Raymond, J., Braddick, O., 1996. Responses to opposed directions of motion: continuum or independent mechanisms? Vis. Res. 36 (13), 1931–1937.
Rodman, H.R., Allbright, T.D., 1989. Single-unit analysis of pattern-motion selective properties in the middle temporal visual area (MT). Exp. Brain Res. 75, 53–64.
Rose, D., Lowe, I., 1982. Dynamics of adaptation to contrast. Perception 11, 505–528.
Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., Tootell, R.B., 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science 268, 889–893.
Simoncelli, E.P., Heeger, D.J., 1996. A model of neuronal responses in visual area MT. Vis. Res. 38 (5), 743–761.
Snowden, R.J., Treue, S., Erickson, R.G., Anderson, R.A., 1991. The response of area MT and V1 neurons to transparent motion. J. Neurosci. 11 (9), 2768–2785.
Solnick, B., Davis, T.L., Sterling, P., 1984. Numbers of specific types of neuron in layer IVab of cat striate cortex. Proc. Natl. Acad. Sci. U. S. A. 81, 3898–3900.
Sutherland, N.S., 1961. Figural aftereffects and apparent size. Q. J. Exp. Psychol. 13, 222–228.
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.
Théoret, H., Kobayashi, M., Ganis, G., Di Capua, P., Pascual-Leone, A., 2002. Repetitive transcranial magnetic stimulation of human area MT/V5 disrupts perception and storage of the motion aftereffect. Neuropsychologia 40, 2280–2287.
Tootell, R.B.H., Reppas, J.B., Dale, A.M., Look, R.B., Sereno, M.I., Malach, R., Brady, T.J., Rosen, B.R., 1995. Visual motion aftereffect in human cortical area MT revealed by functional magnetic resonance imaging. Nature 375, 139–141.
Treue, S., Hol, K., Rauber, H.J., 2000. Seeing multiple directions of motion — physiology and psychophysics. Nat. Neurosci. 3 (3), 270–276.
Van Wezel, R.J.A., Britten, K.H., 2002. Motion adaptation in area MT. J. Neurophysiol. 88, 3469–3476.
Wandell, B., Chial, S., Backus, B., 2000. Visualization and measurement of the cortical surface. J. Cogn. Neurosci. 12 (5), 739–752.
Wohlgemuth, A., 1911. On the aftereffect of seen movement. Br. J. Psychol. 1, 1–117 (Supplement).