Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 45 (2005) 403–412
Vision Research
www.elsevier.com/locate/visres

The motion aftereffect of transparent motion: Two temporal channels account for perceived direction

David Alais [a,*], Frans A.J. Verstraten [b], David C. Burr [c]

[a] Department of Physiology and Institute for Biomedical Research, School of Medical Science, Anderson Stuart Building, University of Sydney, Sydney NSW 2006, Australia
[b] Helmholtz Institute, Psychonomics Division, Utrecht University, Heidelberglaan 2, NL 3584 CS Utrecht, The Netherlands
[c] Istituto di Neurofisiologia del CNR, Via G. Moruzzi 1, Pisa 56010, Italy

Received 23 August 2004

Abstract

Adaptation to orthogonal transparent patterns drifting at the same speed produces a unidirectional motion aftereffect (MAE) whose direction is opposite the average adaptation direction. If the patterns move at different speeds, MAE direction can be predicted by an inverse vector average, using the observer's motion sensitivity to each individual pattern as vector magnitudes. These weights are well approximated by the duration of each pattern's MAE, as measured with static test patterns. However, previous efforts to use the inverse-vector-average rule with dynamic test patterns have failed. Generally, these studies have used spatially and temporally broadband test stimuli. Here, in order to gain insight into the possible contribution of temporal channels, we filtered our test pattern in the temporal domain to produce five ideal, octave-width pass-bands. MAE *durations* were measured for single-component stimuli drifting at various adaptation speeds and tested at a range of temporal frequencies. Then, two components with orthogonal directions and different speeds were combined and the *direction* of the resulting MAE was measured. The key findings are that: (i) for a given adaptation speed, the duration of a single component's MAE is dependent on test temporal frequency; (ii) the direction of MAEs produced by transparent motion (i.e., bivectorial adaptation) also varies strongly as a function test temporal frequency (by up to 90° for some speed pairings); and (iii) the inverse-vector-average rule predicts the direction of the transparent MAE provided the MAE durations used to weight the vector combination were obtained from stimuli matched in adaptation speed and test temporal frequency. These results are discussed in terms of the number and shape of temporal channels in our visual system.

Keywords: Transparent motion; Motion aftereffect; Adaptation; Temporal channels; Speed; Temporal frequency

1. Introduction

The motion aftereffect (MAE) is the term given to the illusory motion perceived in a test pattern following prolonged exposure to a moving pattern (Anstis, Verstraten, & Mather, 1998; Mather, Verstraten, & Anstis, 1998). Typically, the test pattern is stationary and the direction of the illusory movement is opposite that of the adapting movement. Explanations of the MAE are usually couched in terms of adaptation of direction-selective cortical units following a sustained period of activation (Barlow & Hill, 1963; Huk, Ress, & Heeger, 2002; Mather, 1980; Sutherland, 1961). This produces an imbalance in the population of these units such that the direction opposite adaptation is temporarily dominant.

There is a class of motion stimuli that do not produce MAEs in the direction opposite that of adaptation. These are transparent-motion stimuli, produced by superimposing (or rapidly interleaving) two arrays of random dots drifting in differing directions. During

* Corresponding author. Tel.: +612 9351 7615; fax: +612 9351 2058.
E-mail addresses: alaisd@physiol.usyd.edu.au (D. Alais), f.a.j.verstraten@fss.uu.nl (F.A.J. Verstraten), dave@in.cnr.it (D.C. Burr).

0042-6989/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2004.09.005

adaptation, the percept elicited is of two transparent sheets of dots translating independently. As with single-component movement, adaptation to transparent motion also produces a MAE but, curiously, it does not result in two transparent MAEs, one opposite each of the two component motions seen during adaptation. Rather, the MAE of transparent motion is univectorial with a direction opposite the average direction of the adapting motions (Mather, 1980).

Verstraten and colleagues tested the vector-average proposal by adapting to two orthogonal motions and varying their relative speeds. When the adapting components had the same speed, the MAE reflected the average direction. However, when the adapting components differed in speed, the MAE direction deviated from a simple directional average and exhibited a speed-dependent bias. This suggested that the transparent MAE might reflect a true vector average, with the weight of each component in the directional average determined by the motion system's sensitivity to that component's speed. To quantify this, MAE durations were measured for each of the component directions at a range of speeds, and these too were found to depend on adaptation speed. Taking these MAE durations as indicative of the underlying motion sensitivity to a given direction and speed, they can serve as magnitudes for each component's direction, effectively creating motion vectors. For a range of relative speeds, it was found that the speed-dependent deviations from a simple directional average in the transparent MAE were well predicted by an inverse-vector-average rule (Verstraten, Fredericksen, & van De Grind, 1994).

A novel means of eliciting the MAE was introduced first by Mather (1980) and developed by Hiris and Blake (Blake & Hiris, 1993; Hiris & Blake, 1992). Instead of using a static test pattern, they introduced a dynamic test pattern containing an array of dots which jumped about randomly and incoherently. Locally, this stimulus contains a broad range of directions and speeds but (globally) contains no net motion. The rationale for using this stimulus was that a dynamic test stimulus would better drive the motion system during MAE testing to more effectively reveal its adapted state. In a study employing dynamic MAE test stimuli, it was found that MAE directions following transparent motion adaptation differed depending on whether the test stimulus was static or dynamic (Verstraten, van der Smagt, Fredericksen, & van de Grind, 1999). While the MAE directions elicited by static test patterns were predictable using the inverse-vector-average rule (with component MAE durations as weights), those elicited by dynamic test patterns were not (Verstraten and colleagues, unpublished experiments). Their conclusion was that there must be different systems underlying static and dynamic MAEs.

In another study, support was found for two pattern-MAE systems when the durations of MAEs produced by single-component adaptation were measured on static and on dynamic test patterns. The key difference concerned the speed tuning of MAE duration: The strongest dynamic MAEs were elicited by fast adaptation speeds whereas the strongest static MAEs were produced by much slower speeds (Verstraten, van der Smagt, & van de Grind, 1998). Further support for the claim that static and dynamic MAEs are produced in different systems came from a study that showed that it was possible for transparent motion adaptation to produce transparent, bivectorial MAEs (van der Smagt, Verstraten, & van de Grind, 1999). They achieved this by interleaving static and dynamic components in the same test pattern, which presumably tapped separate and independent systems to yield a transparent, bivectorial MAE. Recently, van de Grind and his colleagues collected more evidence for separate mechanisms showing that low and high speeds do not rival when binocularly fused (van de Grind, van Hof, van der Smagt, & Verstraten, 2001).

In the present paper, our goal is to understand why the direction of the dynamically tested transparent MAE cannot be predicted by the inverse-vector-average rule. As a starting point, it is known that dynamic stimuli are processed through temporal channels sensitive to particular ranges of temporal modulation. Psychophysical work on this matter suggests that there are at least two temporal channels (Anderson & Burr, 1985; Fredericksen & Hess, 1998; Hammett & Smith, 1992; Hess & Snowden, 1992; Mandler & Makous, 1984), which can be characterised as a broad low-pass channel and a higher band-pass channel. We reason that adaptation to transparent motion with components of differing speeds would differentially activate the temporal channels. A slow vector would primarily drive the low-frequency temporal channel, whereas a fast vector would primarily drive the high-frequency channel. When the MAE is tested with a broad-band dynamic test pattern, such as those used by Blake and Hiris (1993), the MAE direction would presumably reflect similar contributions from both these channels. However, in this paper, by using dynamic random-dot test patterns temporally filtered into narrow pass-bands, it should be possible to tap preferentially into one temporal channel more than the other. Thus, following adaptation to transparent motion containing both a fast and a slow component, slowly modulating test patterns should elicit a MAE direction opposite the slow component, and quickly modulating test patterns should elicit a MAE direction opposite the fast component. Intermediate test modulations would tap both temporal channels, producing intermediate MAE directions.

The purpose of the present paper is to test these claims and, more specifically, to determine whether the inverse-vector-average rule can accurately predict the direction of transparent MAEs in dynamic test patterns.

2. General methods

2.1. Stimuli and apparatus

Stimuli were generated on a PC using Matlab software in conjunction with a VSG 2/3 card and presented on a Sony GDM-FW900 display running at 144 Hz vertical refresh rate. The stimuli were 256×256 pixel arrays with each element either set randomly to black or white (adaptation patterns) or set to a random luminance value drawn from a flat distribution (test pattern). All arrays were presented at maximum contrast with a 2× magnification in a circular aperture 512 pixels wide containing a fixation point. The aperture subtended 15° of visual angle at the viewing distance of 114 cm. Each (magnified) pixel was square and subtended 0.06° on a side. Average luminance was 42 cd m^{-2}.

2.2. Adaptation pattern

The random binary luminance arrays translated at one of three speeds: (2.3, 9.2, 36.8°/s). The choice of speeds was based on the results of Verstraten et al. (1998). In Experiment 1, a single pattern was displaced, either −45° or +45° upwards from vertical. In Experiment 2, two adapting patterns drifting upward ±45° of vertical were interleaved on alternating frames (72 Hz per pattern) to produce transparent motion. Pixels that fell beyond the aperture were wrapped around and re-entered the aperture at the vertically opposite point.

2.3. Test pattern

A stack of 40 random luminance arrays was generated and passed though a 3D Fourier transformation. In the frequency domain, the 3D array was filtered in the temporal frequency dimension by an ideal, octave-width, band-pass filter which was centred on a frequency of 12.8, 6.4, 3.2, 1.6 or 0.8 Hz. A static test pattern (0 Hz) was also used. The filtered test patterns were therefore spatially broadband with random spatial structure but fairly narrow in temporal frequency. The temporal modulations were produced by cycling through the stack of 40 random luminance images. The images on the display were updated at a rate of 36 Hz (every fourth video frame), producing a cycle length of 1.11 s and a maximum temporal frequency of 18 Hz. Maximum spatial frequency was 8.5 cyc/°.

2.4. Observers

Four observers participated in this study—two authors (DA and FV) and two experienced observers naïve to the aims of the experiments. All observers had normal or corrected-to-normal vision.

3. Experiment 1

We first measured MAE durations for single-direction stimuli for a range of speeds. Three speeds of adapting stimuli were used and the duration of the resulting MAEs was measured with dynamic test stimuli modulating at one of six temporal frequencies. These durations will serve as vector weights to investigate whether the inverse-vector-average rule in Experiment 2 can predict MAE directions for transparent motion adaptation.

3.1. Methods

Observers adapted to univectorial motion for 45 s and indicated the duration of the resulting MAE with a mouse click. There was a pause of not less than 45 s after the MAE ceased. Three adapting speeds and six test temporal frequencies (see Section 2) were combined factorially, making a total of 18 experimental conditions. Adaptation speed, test frequency and motion direction were randomised and counterbalanced. Four trials per condition were run and the MAE durations averaged into a single estimate. Observers were told to note any conditions that failed to produce a detectable MAE. For the purposes of data analysis, these conditions were assigned a value of zero.

3.2. Results and discussion

Fig. 1 shows, for four observers and for three adaptation speeds, MAE duration as a function of test temporal frequency. Two trends are particularly clear: (i) for the fastest adaptation speed (36.8°/s), MAE duration increased with test temporal frequency, and produced no MAE at the lowest test modulation frequency; (ii) for the slowest adaptation speed (2.3°/s), MAE duration decreased with temporal frequency, producing no MAE at the highest test modulation frequency. Durations for the intermediate speed were consistently above zero and exhibited elements of both these trends.

This pattern of results is consistent with the proposal that at least two temporal channels underlie motion perception. The high temporal frequencies produced by the fast adaptation speed adapt the high-frequency channel, and the adapted state of this channel, as measured by MAE duration, is most effectively tapped by high temporal frequency test patterns. By contrast, low temporal test frequencies fail to elicit any MAE following high-speed adaptation. Conversely, the slow adaptation speed adapts the low-frequency temporal channel, as revealed by the strong MAEs for test patterns of low temporal frequency. Low-speed adaptation, however, fails to produce any MAE if the test stimulus contains only high temporal frequencies.

Taken together, these data imply that the fastest and slowest adaptation speeds activate distinct temporal

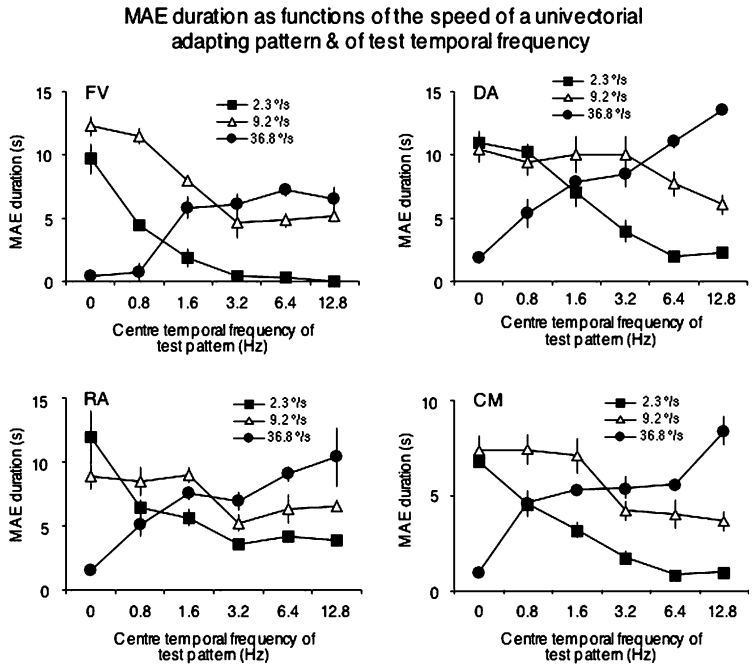

Fig. 1. Results from Experiment 1 for four observers. The data show durations of MAEs produced by univectorial motion adaptation. Three adaptation speeds were examined, and each was tested at six different temporal frequencies. The pattern of results clearly shows that MAE duration depends on both adaptation speed and test temporal frequency. Slow adaptation speeds produce strongest MAEs at low test frequencies, with weak MAEs or none at all, elicited at higher test frequencies. In contrast, fast adaptation speeds produce strongest MAEs at high test frequencies, with weak MAEs or none at all, at higher test frequencies.

channels, and that the range of test frequencies being used is sufficient to tap one or the other channel. Adaptation at the intermediate speed appears to activate both channels since MAEs are evoked by test stimuli at any temporal frequency.

4. Experiment 2

Experiment 2 tests, whether the inverse-vector-average rule can accurately predict the direction of transparent MAEs tested with dynamic test patterns. In this experiment we employ an alternative approach to previous attempts to answer this question by using temporally narrow-band test patterns to probe the motion system after adaptation. In this way, the temporal channels, which presumably underlie the processing of dynamic stimuli, can be studied selectively. If indeed there are only two temporal channels underlying motion processing, narrow-band temporal probes ought to reveal this. Experiment 2 will do this by using transparent motion stimuli whose components have different speeds. For example, we can predict that the MAE direction observed following orthogonal adaptation to a fast and a slow component should depend on test temporal frequency. If a low test frequency is used (low enough to avoid overlap with the higher-frequency temporal channel), MAE direction should be opposite the slow motion vector. If, however, a high test frequency is used (high enough to avoid overlap with the lower channel), MAE direction should be opposite the fast motion vector.

4.1. Methods

The 3 adaptation speeds used in Experiment 1 can be combined into 3 unique speed pairings: 2.3 and 9.2°/s (2-octave speed difference, slow pairing), 9.2 and 36.8°/s (2-octave speed difference, fast pairing), and 2.3 and 36.8°/s (4-octave speed difference). Experiment 2 involved adaptation to these 3 transparent motion stimuli drifting upwards +45° and −45° from vertical, with MAE direction as the dependent variable. Both components of the adapting patterns were visible at all times during adaptation (that is, the univectorial MAEs which resulted were not the result of fused adapting stimuli). The 3 transparent motion stimuli were tested at each of six test temporal frequencies (the same as in Experiment 1, see Section 2), making a total of 18 conditions. Test frequency, and the speed and direction of the vectors were counterbalanced and randomised. On a given trial, adaptation was for 45 s initially, followed by 2 periods of 'top-up' adaptation of 15 s. After each adaptation period, observers indicated MAE direction using a mouse-operated rotating pointer. These 3 MAE settings were averaged into a single datum for that trial. Five trials per condition were run and averaged into a single estimate of MAE direction for each condition. There was an interval of not less than 60 s between

trials. Observers were asked to note all conditions that did not yield a discernible MAE, although none was observed.

4.2. Results and discussion

Fig. 2 shows, for four observers and for the 3 component-speed pairings, MAE direction as a function of test temporal frequency. Three trends are evident in all observers. The strongest of these occurs with the 4-octave speed pairing (2.3 and 36.8°/s) where perceived MAE direction ranges through approximately 90°, from opposite the slow vector at low test temporal frequencies, to opposite the fast vector at high test frequencies. This confirms the prediction made in the introduction to this experiment that only high-speed adaptation would be tapped by the high temporal test pattern, and only low-speed adaptation would be tapped by the low frequency pattern.

Both of the 2-octave speed pairings yielded directional shifts that were equivalent to each other in magnitude. However, both spanned roughly half the directional range – approximately 45° – observed for the 4-octave speed difference. For the slower of the 2-octave pairings (2.3 and 9.2°/s), MAE directions reflected a simple average direction when test temporal frequencies were low, and shifted progressively to opposite the faster speed as test frequency increased. This is exactly the pattern expected given the data of Experiment 1, which showed that the slow speed drives only the low-frequency temporal channel whereas the intermediate speed drives both channels. Thus, any test frequency that drives only the low temporal frequency channel would produce MAEs opposite the average direction, as this channel would be adapted by both directions. High test frequencies, however, should elicit a MAE direction opposite the faster speed, as this speed would adapt only the high-frequency channel. Conversely, for the faster of the 2-octave pairings (9.2 and 36.8°/s), MAE directions ranged from opposite the intermediate speed when test frequencies were low to a simple average direction when test temporal frequencies were high. Again this is expected as the low-frequency channel only responds to the intermediate speed while the high-frequency channels respond to both speeds. These patterns of results have important implications for the shapes of the underlying temporal channels and will be discussed further in Section 5.

Can these MAE directions of transparent motion be accounted for by an inverse-vector-average rule? This

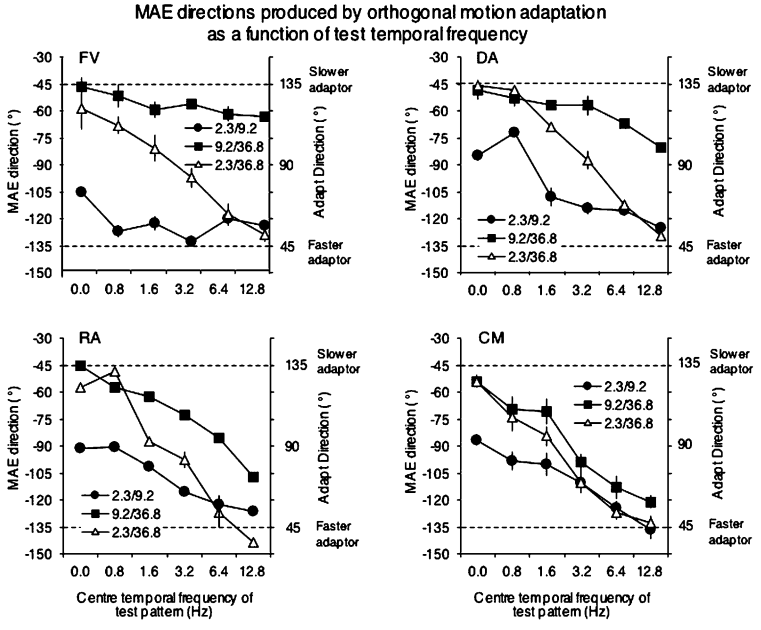

Fig. 2. Results from Experiment 2 for four observers shown on double *Y*-axis plots showing MAE direction (left axis) and adaptation direction (right axis). As in Experiment 1, the 4-octave speed difference makes the clearest prediction, and the data bear it out: slow test temporal frequencies resulted in MAE directions opposite the slow adapting component (2.3°/s) and fast test temporal frequencies result in MAEs opposite the fast adapting component. The simplest way to account for an effect of such a monotonic change in direction from opposite one component direction to opposite the other is by the trading-off of two overlapping temporal channels. That is, there is no point at which no MAE results, and a smooth change in direction results from one extreme to the other. For the other two speed pairs, the MAE directions move monotonically from one extreme to the middle and no further. This means, to take the slow speed-pair (2.3 and 9.2°/s), that directions go smoothly from opposite the slow component to the average direction, and for the fast speed-pair (9.2 and 36.8°/s) that directions move from opposite the fast vector to the average direction. This is again consistent with a trading-off of two temporal channels which goes smoothly and progressively. In short, it shows that the dynamic MAE is not the result of a single dynamic system. Instead, it points to two sub-systems underlying the dynamic MAE which most likely correspond to two temporal channels, as described in the literature.

can be tested by taking from Experiment 1 the appropriate MAE durations and calculating, for each transparent motion condition in Experiment 2, the inverse tangent to obtain their predicted vector-average direction. Ideally, MAE velocities from Experiment 1 would be used to make these predictions, however, duration is a measure that correlates well with velocity as an index of the MAE (Keck, Palella, & Pantle, 1976). Predicted MAE directions calculated from the Experiment 1 (duration data) are plotted in Fig. 3 for the 4 subjects. In all cases, the predictions capture the essential characteristics of the data from Experiment 2 (Fig. 2). That is, the full 90° range in MAE direction is predicted for the 4-octave speed pair, while predictions for the 2-octave speed pairs cover half that range, with each pair occupying different halves of the MAE direction range. The determinants of this particular pattern of results are discussed in Section 5.

5. General discussion

One interesting aspect of these data is that a MAE is not an inevitable consequence of motion adaptation. If low speeds are used to adapt, then no MAE will be observed if a high temporal frequency test stimulus is used. The inverse also applies: Adapting to a high-speed pattern will not produce a MAE if a low-frequency test pattern is used. The data from Experiment 1 (Fig. 1) illustrate this point. A second interesting aspect of these data is that there is no single MAE direction predetermined by a given pair of adaptation directions. The directions contained in the adapting stimuli merely constrain a range of possible outcomes. The specific MAE direction that obtains is dependent on the test temporal frequency (see Fig. 2). It is not possible to account for these findings with a single temporal frequency channel.

5.1. The shape of the temporal channels

Earlier studies using adaptation or masking (Anderson & Burr, 1985; Fredericksen & Hess, 1998; Hammett & Smith, 1992) pointed to the existence of two temporal channels of quite different shapes, illustrated schematically in Fig. 4(c). While the fast channel was characterised as band-pass, the slow channel was low-pass and very broad, extending almost as far into the high temporal frequency range as the fast channel. A simple imple-

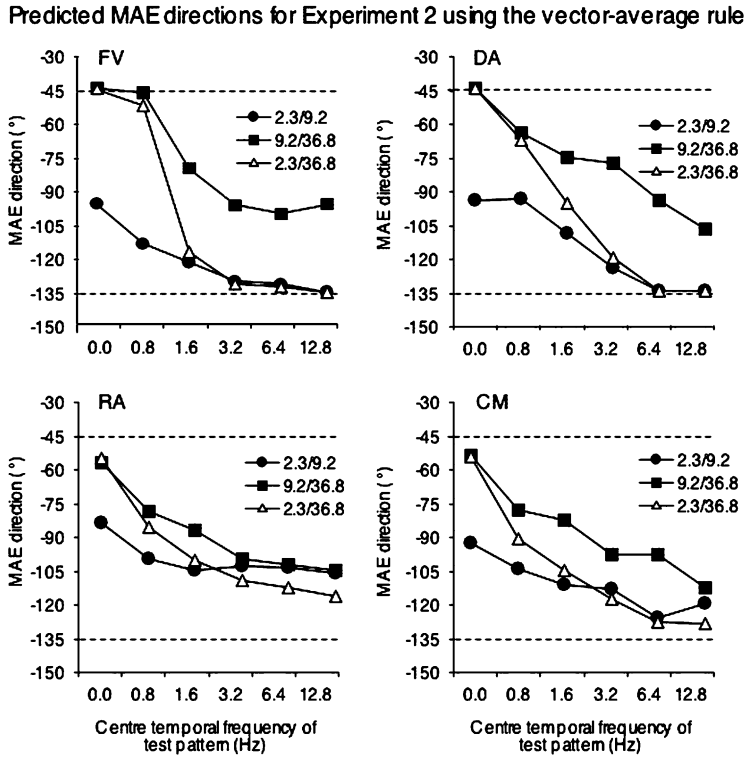

Fig. 3. Predicted MAE directions calculated using the inverse-vector-average rule. The vector magnitudes are taken from Experiment 1 (see Fig. 1) in which MAE durations for single-vector stimuli were measured. The essential characteristics of the bivectorial MAE data presented in Fig. 2 are captured by the inverse-vector-average predictions. Specifically, as test temporal frequency increases: (i) MAE direction for the 4-octave speed pairing ranges from opposite the slow to opposite the fast adaptation direction; (ii) the slower of the 2-octave speed pairings ranges from opposite the average direction to opposite the faster direction; (iii) the faster of the 2-octave speed pairings ranges from opposite the slower direction to opposite the average direction. Note that the test frequencies in Experiments 1 and 2 are the same and that the bivectorial stimuli in the second experiment are simply pairings of the same stimuli as used in the first experiment.

mentation of this model predicts an asymmetry in MAE direction that was not observed in our data. There is also some evidence to support the existence of a second band-pass filter in a three-channel model (Hess & Snowden, 1992; Johnston & Clifford, 1995; Mandler & Makous, 1984), although the existence of a broad low-pass channel in this scheme would still predict asymmetrical MAE directions. The two-channel model shown in Fig. 4(c) predicts that low temporal test frequencies should always elicit a MAE, regardless of the speed (and temporal content) of adaptation, because the low-frequency channel spans virtually all adaptation speeds. Because of this, a test pattern modulating at any frequency would suffice to drive it and reveal its adapted state. In contrast, the data from Experiment 1 show that low temporal test frequencies only elicit a MAE following low-speed adaptation (Fig. 1). The broad low-pass model also predicts that high temporal test frequencies should elicit MAEs following adaptation at almost any speed, again due to the very broad tuning of the low-pass channel. Our data show instead that high temporal test frequencies only elicited MAEs following adaptation to high speeds (Fig. 1).

We can conclude from Fig. 1 that the range of temporal frequencies used in these experiments spanned more than the overlapping portion of two temporal channels, so that the extremes of our range tapped distinct populations of temporally selective neurons. This suggests a different scheme from that implied by the broad low-pass model. More specifically, the complementary symmetry of the slow- and fast-adaptation curves in Fig. 1 is telling as it implies two partly overlapping band-pass channels. This arrangement, shown in Fig. 4(a), would produce the complete trade-off between low- and high-speed adaptation seen in Fig. 1; the broad low-pass model (Fig. 4(c)) would not. Moreover, the overlapping band-pass model can also predict the transparent MAE directions obtained in Experiment 2. Fig. 4(c) shows predictions based on a model of two, overlapping, Gaussian-shaped temporal channels illustrated in Fig.

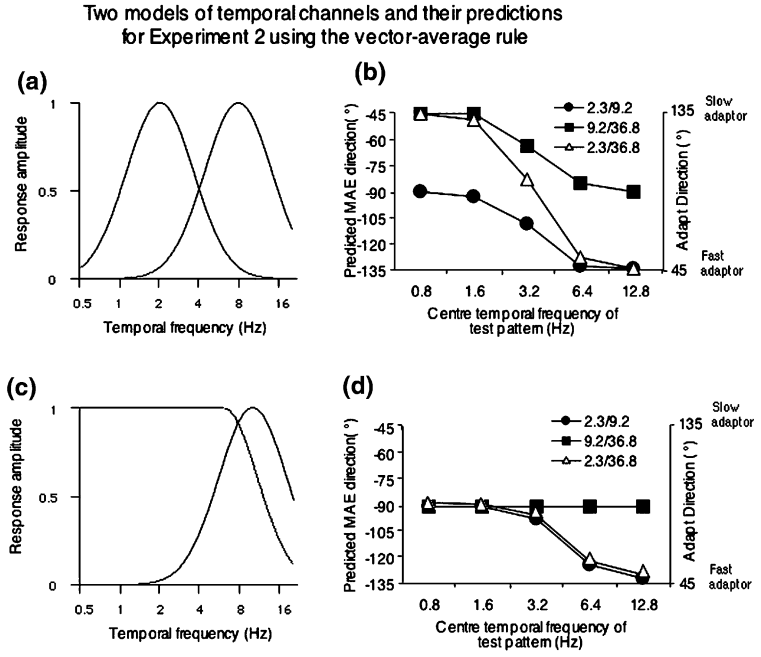

Fig. 4. (a) A depiction of the overlapping band-pass model proposed in the discussion. Both channels are Gaussian on a log temporal frequency axis. The two channels shown here have a full bandwidth of 2 octaves and are centred at 2 and 8 Hz. (b) Predictions of bivectorial MAE direction as a function of test temporal frequency based on the inverse-vector-average rule and outputs from a simple implementation of the overlapping band-pass model (panel A, this figure). The test frequencies and adapting directions and speeds are the same as those used in Experiment 2 (see Fig. 2). The resulting pattern of MAE directions matches the characteristics of the empirical data (Fig. 2). (c) A generic version of the 'broad low-pass model' proposed by Anderson and Burr (1985), Hammett and Smith (1992), and Fredericksen and Hess (1998), with its high-frequency channel centred at 10 Hz with a full bandwidth of 2 octaves and a broad low-pass channel extending almost as far as the high-frequency channel. (d) Predictions of bivectorial MAE direction as a function of test temporal frequency based on the inverse-vector-average rule and outputs from the generic broad low-pass model (panel C, this figure). Again, frequencies, directions and speeds are as for Experiment 2. The resulting pattern of MAE directions is not at all like the empirical data obtained in Experiment 2 (Fig. 2). In particular, the model fails to predict the MAE direction opposite the slower adaptor that was observed empirically at low test temporal frequences (Fig. 2). This is due to the breadth of the low-pass channel: both adapting directions, whether fast or slow, will adapt the low-pass channel and result in an MAE (for that channel) opposite the average adaptation direction. If one of the adaptation speeds is fast enough to drive the high-frequency channel, this will shift the MAE towards a direction opposite the faster adaptor. However, there is no speed/frequency combination that can shift the 'average direction' being signalled by the broad low-pass channel towards the slower adaptation direction.

4(a). For the 4-octave speed pair (comprising a low and a fast speed), MAE direction moves progressively from opposite the slow component to opposite the fast component as test temporal frequency increases. This is because the adapting speeds drive separate temporal channels and because the extremes of the test frequency range tap these channels uniquely. For the low-speed 2-octave pair (comprising a low and a moderate speed), low test frequencies produce a simple average direction, since both speeds adapt the low frequency channel. For high test frequencies, only the high-frequency channel is tapped, and so only the higher of the two speeds (which in fact adapted both channels) is able to contribute to the MAE, yielding a direction directly opposite the faster vector. The same logic explains why MAE directions for the high-speed 2-octave pair (a moderate and a high speed) ranged from opposite the slow speed at low test frequencies to a simple direction average at high test frequencies.

Fig. 4(d) shows predictions from a simple implementation of the broad-low-pass model. Because of the model's asymmetry in having such a broad low-pass channel, it cannot account for the full range of MAE directions plotted in Fig. 2. It predicts a simple directional average when lower test frequencies are used (these probe only the lower frequency channel which, being broad, has adapted to both the fast and slow directions) and a direction opposite the fast component for high test frequencies. This model is unable to explain our finding of MAE directions opposite the slow vector at low test frequencies because its low-frequency channel can never be driven solely by low-frequency input. In our conditions, where the transparent adapting stimuli always had components of different speeds, the broad low-frequency channel would always respond to both components. In contrast, a more symmetric model such as the one illustrated in Fig. 4(a), will predict the symmetry of the data: That adapting to high temporal frequency produces little MAE at low temporal frequency, and vice versa.

In sum, it seems clear that two overlapping band-pass channels could account well for both the duration data (Experiment 1) and the direction data (Experiment 2) [1]. However, some caveats are necessary. Importantly, we cannot estimate the bandwidth of the channels or their centre frequencies, as the adapting stimuli were spatially broadband, so that for a given adaptation speed, adaptation at a broad range of temporal frequencies took place (it is noteworthy, though, that at 8.5 cyc/°, the maximum spatial frequency was not especially high). Future experiments, using spatially band-limited stimuli to accurately control temporal frequency, are already underway to clarify the width and centre frequencies of these temporal channels. Given this limitation, the specific shape, location and width of the channels shown in Fig. 4a are included only for illustrative purposes (while their implied shape is bandpass, the channels need not be strictly Gaussian, see Johnston & Clifford (1995) and Hess & Snowden (1992)). They do suffice, however, to illustrate the principle that a minimum of two channels, arranged so that a given temporal frequency will always drive one filter more than the other, can produce a pattern qualitatively similar to the data we obtained in Experiment 2 (cf. Figs. 2 and 4(b)). For the sake of illustration, the model in Fig. 4(a) shows two channels centred at 2 and 8 Hz, each with a full bandwidth of 2 octaves.

Although a model of two overlapping band-pass channels appears to explain the symmetry of our data well, existing broad low-pass models could also account for the data if suitably extended. Much of the evidence supporting existing temporal channel models comes from masking and threshold experiments that are likely to reflect the temporal characteristics of early detectors. The perception of motion presumably arises at a later stage and intervening operations may take place on the outputs of these early mechanisms before motion is computed. The overlapping band-pass scheme (Fig. 4(a)) would not need such operations because motion could be computed simply by the ratio of high to low frequency outputs (i.e., high bandpass/low bandpass). However, for the broad low-pass model shown in Fig. 4(c), if the low-pass channel were first divided by the band-pass channel, producing an output function resembling the low band-pass channel in Fig. 4(a), motion could then be computed by dividing this result into the band-pass channel (i.e., bandpass/[lowpass/bandpass]). While these operations are purely hypothetical, it does serve to demonstrate that our MAE data are not necessarily incompatible with standard temporal frequency models.

Whether or not such operations take place cannot be determined on the basis of the present data, however, it is interesting to note that there is plenty of scope for such operations to occur. The reason is that MAEs presumably require motion opponency, which is observed in area MT but not in primary visual cortex (Heeger, Boynton, Demb, Seidemann, & Newsome, 1999). Also, our motion stimulus was a transparent motion display, and this form of motion is reliably signalled in MT but not in V1 (Qian & Andersen, 1994, 1995). Thus, one way to resolve the differences between the traditional broad low-pass model and the temporal filter model implied by our MAE data is in terms the stages of visual processing being tapped by the measurement techniques. Masking and detection measurements probably reflect the low-level temporal characteristics of

[1] For simplicity, the focus to this point has been on temporal filtering in just two channels. To explain direction, a bank of temporal filters with signed outputs specifying a particular orientation would be needed.

units in V1, while MAE data necessarily reflect a motion opponency stage not observed prior to MT, leaving.

If this difference in the level of processing being tapped were to explain the disagreement between our results and those of prior temporal filtering studies, then using our temporally filtered stimuli in a masking or detection paradigm ought to produce results matching those of traditional temporal filter models. However, if this hypothetical experiment were conducted and the implied filters matched those suggested by our current MAE measurements, they would point to another interesting possibility: That in using a temporal stimulus with random spatial structure, our stimuli only activate the dynamic (motion system). Previous temporal channels studies have used gratings or Gabors which would likely have activated units in the form pathway (at least when undergoing low frequency temporal modulations). One consequence of this could be to obscure the roll-off in sensitivity at low frequencies in the motion system suggested by Fig. 4(a), producing a pattern of data resembling a broad low-pass channel but which could actually be a summation of a low band-pass channel (Fig. 4(a)) with a form channel maximally sensitive at 0 Hz and decaying to zero at a relatively low frequency. If these channels crossed at half-height, their sum would resemble the broad low-pass channel shown in Fig. 4(c). Future experiments employing our temporal filtered, spatially random stimulus in low-level detection experiments should reveal which of these alternatives is correct.

6. The inverse-vector-average rule

One of the aims of this paper was to determine whether the inverse-vector-average rule, implemented with univectorial MAE durations as vector weights, accurately predicts the direction of transparent MAEs for dynamic test patterns. The agreement between the main features of the data plotted in Fig. 2 and the predictions plotted in Fig. 3 indicate that it does. Presumably, it does so because the MAE durations produced by adaptation to a single motion are reasonable indicators of the visual system's sensitivity to that motion. Motions to which the system is highly sensitive will provoke a rigorous neural response, which will in turn produce strong adaptation. Conversely, motions engendering only weak neural responses will lead to weak adaptation. For this reason, using the inverse-vector-average rule with univectorial MAE durations as weights is just a small step removed from using actual motion sensitivities to weight the combination. Explicated like this, the logic of the inverse-vector-average rule becomes clear, as does its primary advantage: Transparent MAE directions can be accounted for without the lengthy threshold experiments needed to quantify motion sensitivity.

If there is a disadvantage to using univectorial MAE durations to weight the inverse-vector-average rule, it is that judging MAE cessation with a stable criterion for MAEs varying widely in strength is not easy. Nor is it easy to decide if a MAE is non-existent because the measurement requires a certain inspection time. We dealt with this problem by training observers with a range of MAE strengths to help them establish a stable criterion, and by instructing them to indicate 'no aftereffect' for conditions in which, after reasonable inspection, they were sure that no MAE had occurred. In these cases, a value of zero was substituted for the spurious duration recorded from the belated mouse-click. With these basic controls in place we were able to predict transparent MAE directions that accord well with the features of the actual data. Perhaps for effects not expected to vary as broadly as the 90° range in MAE direction we report, computing motion sensitivities from threshold experiments might serve better as vector weights since threshold measurements are probably more robust and stable than MAE-duration data. However, both methods have been compared directly using static test patterns and MAE directions were roughly equally well predicted using either value to weight the directional combination (Pantle, 1998; Verstraten, 1994; Verstraten et al., 1994).

With the inverse-vector-average rule found to work so effectively in Experiment 2, what accounts for the prior failures to extend this rule from statically tested MAEs to dynamically tested MAEs? These may be attributed to two main shortcomings. First, although the possibility that static and dynamic MAEs might depend on different cortical systems was discussed by Verstraten et al. (1999), due consideration was not given to the evidence indicating that the dynamic system itself could be further divided into two or possibly three temporal channels. Second, the use of temporally broadband test patterns would have provided strong activation of both temporal channels, tending to obscure any distinction between them. This would explain why transparent MAE directions tested with broadband dynamic stimuli could not be predicted by using static MAE durations as vector weights, or combinations of static and dynamic MAE durations as weights (van der Smagt, Personal communication). To illustrate the second point, consider a slow adapting pattern which would create adaptation only in the low-frequency temporal channel. Probing for evidence of this adaptation with a broad-band temporal pattern would activate both the low- and the high-frequency temporal channels with noise. This would swamp the adaptation 'signal' with additional spurious activity from the high-temporal frequency channel, effectively diluting the strength of the MAE and blurring any distinction between the channels.

To summarise, the inverse-vector-average rule does successfully predict the direction of transparent MAEs

tested with dynamic test patterns, provided that the strengths of the component MAEs are measured with narrow temporal probes. Use of temporally broad-band test probes will compromise predicted MAE directions in two ways: (i) it constrains the range of transparent MAE directions because testing will always strongly activate both channels during testing, limiting deviations from the average direction; (ii) it reduces the duration of the univectorial duration estimates which serve as vector weights, because activating both channels during testing will dilute the adapted channel with a second, unadapted channel. The first point may explain why several reports noted that MAEs tested with (broadband) dynamic stimuli tend to be shorter than MAEs tested with static test patterns (Green, Chilcoat, & Stromeyer, 1983; Verstraten et al., 1994).

In a final and general observation, we note that while spatial filtering of broadband stimuli has long been a standard approach in vision research, temporal filtering of broadband stimuli has largely been ignored. The work presented here is a rare example in vision research of filtering in the temporal domain. From our current findings we can see the power of a temporal-filtering approach for studying temporal aspects of vision.

Acknowledgements

D.A. was supported by a Marie Curie Fellowship from the European Commission and FV by the Netherlands Organisation for Scientific Research (NWO-PIONIER). We thank Maarten van der Smagt for helpful discussion.

References

Anderson, S. J., & Burr, D. C. (1985). Spatial and temporal selectivity of the human motion detection system. *Vision Research, 25*, 1147–1154.

Anstis, S. M., Verstraten, F. A. J., & Mather, G. (1998). The motion aftereffect. *Trends in Cognitive Science, 2*, 111–117.

Barlow, H. B., & Hill, R. M. (1963). Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects. *Nature, 200*, 1345–1347.

Blake, R., & Hiris, E. (1993). Another means for measuring the motion aftereffect. *Vision Research, 33*, 1589–1592.

Fredericksen, R. E., & Hess, R. F. (1998). Estimating multiple temporal mechanisms in human vision. *Vision Research, 38*, 1023–1040.

Green, M., Chilcoat, M., & Stromeyer, C. F. III, (1983). Rapid motion aftereffect seen within uniform flickering test fields. *Nature, 304*, 61–72.

Hammett, S. T., & Smith, A. T. (1992). 2 temporal channels or 3 – a reevaluation. *Vision Research, 32*, 285–291.

Heeger, D. J., Boynton, G. M., Demb, J. B., Seidemann, E., & Newsome, W. T. (1999). Motion opponency in visual cortex. *Journal of Neuroscience, 19*, 7162–7174.

Hess, R. F., & Snowden, R. J. (1992). Temporal properties of human visual filters: Number, shapes and spatial covariation. *Vision Research, 35*, 47–59.

Hiris, E., & Blake, R. (1992). Another perspective on the visual motion aftereffect. *Proceedings of the National Academy of Sciences of the United States of America, 89*, 9025–9028.

Huk, A. C., Ress, D., & Heeger, D. J. (2002). Neuronal basis of the motion aftereffect reconsidered. *Neuron, 32*, 161–172.

Johnston, A., & Clifford, C. W. (1995). A unified account of three apparent motion illusions. *Vision Research, 35*, 1109–1123.

Keck, M. J., Palella, T. D., & Pantle, A. (1976). Motion aftereffect as a function of the contrast of sinusoidal gratings. *Vision Research, 16*, 187–191.

Mandler, M. B., & Makous, W. (1984). A three channel model of temporal frequency perception. *Vision Research, 24*, 1881–1887.

Mather, G. (1980). The movement aftereffect and a distribution–shift model for coding the direction of visual movement. *Perception, 9*, 379–392.

Mather, G., Verstraten, F. A. J., & Anstis, S. (1998). *The motion aftereffect: A modern perspective*. Cambridge, MA: MIT Press.

Pantle, A. (1998). How do measures of motion aftereffect measure up? In G. Mather, F. A. J. Verstraten, & S. M. Anstis (Eds.), *The motion aftereffect: A modern perspective*. Cambridge: MIT Press.

Qian, N., & Andersen, R. A. (1994). Transparent motion perception as detection of unbalanced motion signals. II. Physiology. *Journal of Neuroscience, 14*, 7367–7380.

Qian, N., & Andersen, R. A. (1995). V1 responses to transparent and nontransparent motions. *Experimental Brain Research, 103*, 41–50.

Sutherland, N. S. (1961). Figural aftereffects and apparent size. *Quarterly Journal of Experimental Psychology, 13*, 222–228.

van de Grind, W. A., van Hof, P., van der Smagt, M. J., & Verstraten, F. A. J. (2001). Slow and fast visual motion channels have independent binocular rivalry stages. *Proceedings of the Royal Society of London B, 268*, 437–443.

van der Smagt, M. J., Verstraten, F. A. J., & van de Grind, W. A. (1999). A new transparent motion aftereffect. *Nature Neuroscience, 2*, 595–596.

Verstraten, F. A. J. (1994). Multi-vectorial motion and its after-effect(s). Unpublished doctoral dissertation. Utrecht: Utrecht University.

Verstraten, F. A. J., Fredericksen, R. E., & van De Grind, W. A. (1994). Movement aftereffect of bi-vectorial transparent motion. *Vision Research, 34*, 349–358.

Verstraten, F. A. J., van der Smagt, M. J., Fredericksen, R. E., & van de Grind, W. A. (1999). Integration after adaptation to transparent motion: Static and dynamic test patterns result in different aftereffect directions. *Vision Research, 39*, 803–810.

Verstraten, F. A. J., van der Smagt, M. J., & van de Grind, W. A. (1998). Aftereffect of high-speed motion. *Perception, 27*, 1055–1066.