

Chapter 5

Higher-Order Effects

*Jody Culham, Shin'ya Nishida, Timothy Ledgeway,
Patrick Cavanagh, Michael von Grünau, Michelle
Kwas, David Alais, and Jane Raymond*

In this chapter we consider some potential high-level influences on the strength and direction of the motion aftereffect (MAE), and their implications for psychophysical and neural theories of motion perception. We focus on a wide range of stimuli and paradigms, including dynamic tests, higher-order stimuli, multivectorial motion, and brief adaptation periods. We also consider the effects of the observer's state of attention and the demands of the task, showing that both can influence the direction and size of the effect. These issues suggest that the MAE is not necessarily a unified phenomenon resulting only from low-level, preattentive processing of unidirectional motion from luminance cues. Rather, the MAE may have a wide range of influences and consequences, some of which are dependent on the particular stimuli and the observer's state of mind.

5.1 Static vs. Dynamic Motion Aftereffect Tests

Prolonged exposure to moving stimuli generates illusory backward motion in test stimuli presented subsequently at the same location. This MAE conventionally has been studied using static test patterns, although recent studies have also employed dynamic test patterns such as counter-phasing gratings and random dots which are directionally ambiguous (figure 5.1). Dynamic test stimuli were originally used for a number of technical reasons. For instance, directional components within the stimuli are easy to manipulate, a useful property when MAE magnitude is evaluated in terms of the motion signals required to cancel the aftereffect (von Grünau, 1986; Blake and Hiris, 1993; Ledgeway, 1994). The illusory opposite motion seen in dynamic tests is perceptually indistinguishable from real motion (Levinson and Sekuler, 1975; Hiris and Blake, 1992), so it is easier for subjects to judge than that seen in static tests. In addition, it was suggested that perceptually unstable test stimuli might provide a more sensitive technique for revealing weak adaptation effects (von Grünau and Dubé, 1992).

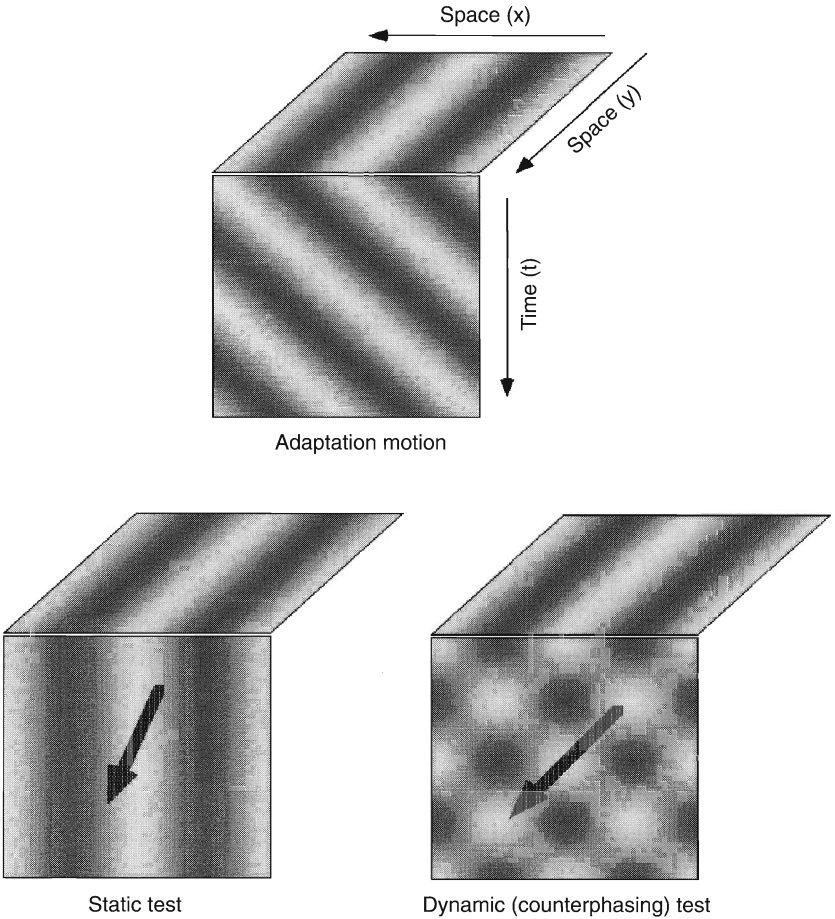


Figure 5.1

Motion aftereffects measured using static and dynamic (counterphase) gratings. Top: Adaptation stimulus. A rightward drifting sinusoidal grating is shown by a space (x)–space (y) and space (x)–time (t) plot. Bottom left: Static test grating. Arrow indicates the direction of the aftereffect. Bottom right: Counterphase test grating.

It has been suggested that the MAE is a result of selective adaptation of directionally selective mechanisms that respond continuously to the adaptation motion. A parsimonious explanation of the induction of MAEs either in static or dynamic test stimuli is that they both result from adaptation at the same site of visual motion processing. There appears to be no a priori reason to doubt this assumption, and it is implicit in several studies that have used dynamic test stimuli mainly for technical reasons. It has recently been found, however, that the MAE measured using dynamic test stimuli (dynamic MAE) demonstrates a number of differences from the MAE measured using static test patterns (static MAE). This is hard to reconcile with the proposition that a single mechanism is responsible for both types of MAE. A marked difference concerns the effects of second-order stimuli and higher-order motion, which are discussed later in this chapter. In this section, we first compare the properties of static and dynamic MAEs revealed by studies that have mainly employed first-order (luminance-defined) motion stimuli.

5.1.1 Contrast Dependency

Luminance contrast is one of the stimulus parameters that strongly control the magnitude of the classic static MAE (Keck et al., 1976). The effects of contrast on the duration of the dynamic MAE (seen in a 2.0-Hz counterphasing grating presented in the central visual field) were found to be quite similar to those known for the static MAE (Nishida et al., 1997a). In each case, as the adaptation contrast is increased, MAE duration steadily increases, leveling off when adaptation contrast approaches test contrast. As test contrast is increased, the MAE duration decreases, with a slightly steeper slope in the case of the dynamic MAE. These results appear to favor the suggestion that a common process is responsible for both static and dynamic MAEs, although the similarity could, in principle, arise from analogous adaptation processes occurring at different sites of visual motion processing. Another notable point with regard to the effects of contrast is that adaptation contrasts as low as two times the detection threshold can induce reliable MAEs in either type of test stimulus (Nishida et al., 1997a). This suggests that such differences that are present in sensitivity to the two types of test stimulus are moderate, at least in the case of simple luminance gratings.

5.1.2 Spatial Frequency

Static MAE magnitude measured with luminance gratings is typically strongest when the adaptation and test stimuli are matched for spatial frequency (Over, Broerse, et al., 1973; Cameron et al., 1992). This implies that the static MAE is spatial frequency-selective, and therefore that the underlying mechanisms responsible for this MAE are each narrowly tuned

to a limited range of spatial frequencies. Ashida and Osaka (1994) and Bex et al. (1996) investigated whether such selectivity is also obtained with dynamic MAEs. Ashida and Osaka found that the duration of MAE measured with sinusoidal counterphasing gratings was affected by the spatial frequencies of the adaptation and test stimuli. Nevertheless, when their data are plotted as functions of the adaptation spatial frequency, the maximum MAE generally occurs at a frequency lower than the test frequency. On the other hand, Bex et al. found clear spatial frequency selectivity of the MAE even though they also used counterphasing tests. One important procedural difference between these two studies was the test temporal frequency: Ashida and Osaka used 5.0 Hz while Bex et al. used 0.25 Hz. A follow-up study (Mareschal, Ashida, et al. 1997) found that the spatial frequency selectivity of the MAE broadened as the test temporal frequency increased, and almost disappeared at temporal frequencies equal to 2 Hz. These results indicate broad spatial-frequency tuning of the underlying mechanisms that give rise to MAEs with counterphasing gratings equal to 2 Hz. Von Grünau and Dubé (1992), however, using a nulling method and plaid adaptation stimuli, reported results that appear to indicate spatial-frequency selectivity of MAE seen in a 2-Hz counterphasing test.

5.1.3 Location Specificity

The static MAE is typically observed only in the region directly stimulated during adaptation (Anstis and Gregory, 1965). However, the dynamic MAE does not show such retinotopic specificity. Von Grünau and Dubé (1992) measured the magnitude of the dynamic MAE for tests presented at the adapted location and at another location equidistant from fixation. They reported dynamic MAEs at both locations, though the remote MAE was weaker and showed larger intersubject variability. Furthermore, MAEs at the test location were tuned for the direction and spatial-frequency of the adapting grating or plaid; however, remote MAEs showed no such direction tuning (responding with approximately the same magnitude even to gratings differing in orientation and direction by up to 60 degrees) nor spatial-frequency tuning.

5.1.4 Temporal Frequency

As the speed or the temporal frequency of the adaptation stimulus increases, the magnitude of the MAE first increases and then decreases. When this inverse-U function is measured for various spatial frequencies, it is possible to tell whether the speed or the temporal frequency characterizes this effect. Previous studies using static test stimuli (Pantle, 1974; Wright and Johnston, 1985) showed that MAE magnitude depends on adaptation temporal frequency. That is, MAE magnitude generally peaks

at adaptation temporal frequencies of 5 to 10 Hz regardless of the adaptation spatial frequency. On the other hand, Ashida and Osaka (1995b) reported that the duration of the dynamic MAE (measured using 2.5- and 5.0-Hz counterphasing gratings) depends on the adaptation speed, rather than temporal frequency, and peaks at adaptation speeds of 5 to 10 degrees/second irrespective of the adaptation spatial frequency. These results suggest that the mechanism underlying the static MAE codes temporal changes in luminance in terms of temporal frequency, while that underlying the dynamic MAE codes speed.

In the case of the dynamic MAE, the effects of test temporal frequency have also been examined. Bex et al. (1996) found a decrease in MAE duration with increasing test temporal frequency (except when the test was nearly static; i.e., at very low temporal frequencies). The shape of this function was not affected by the spatiotemporal frequency of the adaptation stimulus and thus the effect was selective neither for speed nor for temporal frequency. On the other hand, Verstraten, van Wezel, et al. (1994c) reported results that suggest an interaction between the temporal frequencies of the adaptation and test stimuli. They measured the direction of the MAE using dynamic random-dot tests after adaptation to a transparent motion stimulus in which two random-dot patterns drifted in orthogonal directions at different speeds. Their results suggest that the influence of the slower components in the adaptation stimulus is reduced as the test temporal frequency is increased.

5.1.5 Interocular Transfer

The magnitude of interocular transfer (the extent to which adaptation presented to one eye produces aftereffects measured with the other eye; see chapter 4) has been used as a tool to investigate the binocularity of the mechanisms responsible for aftereffects. It is known that the static MAE transfers interocularly, although the magnitudes of the MAE measured with the nonadapted eye are about half those measured with the adapted eye (Wohlgemuth, 1911; Moulden, 1980; see chapter 4). This result, called partial transfer, indicates that the mechanism subserving the static MAE is neither totally monocular nor binocular. In comparison, Nishida et al. (1994) found that transfer of the dynamic MAE, measured with a 2-Hz counterphasing grating presented in the central visual field, is nearly perfect regardless of the type of adaptation stimuli. This result led them to conclude that the dynamic MAE reflects adaptation at a site where binocular signals are completely integrated and eye-of-origin information is lost.

However, interocular transfer is not perfect for all MAEs measured using dynamic tests. Green, Chilcoat, et al. (1983) found no interocular transfer for the MAE that is seen in a homogeneous rapidly flickering test

field after adaptation to a drifting grating of low spatial frequency. On the other hand, Steiner et al. (1994) found interocular transfer using a different type of dynamic test stimulus composed of random dots moving incoherently in all directions, but the magnitude of transfer was significantly less than perfect (76 percent on average) after adaptation to translational motion. The lack of interocular transfer reported by Green et al. (1983) is incompatible with the results of all other studies. It appears that the MAE seen in patterned dynamic stimuli and that seen in patternless flicker are different phenomena. This idea is strengthened by the finding that second-order motion is differentially effective in producing these two MAEs (Nishida and Sato, 1995). As for the discrepancy between Nishida et al. (1994) and Steiner et al. (1994), a notable difference is the method used to evaluate the magnitude of MAEs: Nishida et al. measured the MAE duration, while Steiner et al. measured the amount of motion signal required to null the adaptation effects.

To test whether different methods could give rise to different magnitudes of interocular transfer, Ashida and Nishida (unpublished) compared the MAE duration and the nulling motion strength obtained under the same stimulus conditions (adapting with a drifting grating, and testing with a counterphase grating). Their preliminary results suggest that interocular transfer is nearly perfect when measured by MAE duration, but less than perfect when measured by nulling motion strength. Another factor that may be relevant to the discrepancy between studies is the temporal structure of the test stimulus. While the temporal frequency of a counterphasing grating is localized at the modulation frequency, incoherently moving dots contain a broad range of temporal frequencies, including static and very slowly drifting components. Transfer (based on MAE duration measured with counterphasing gratings) is perfect only when the modulation rate is higher than 1 Hz (Ashida, Verstraten, et al., 1997); thus slow components in the test stimulus of Steiner et al. (1994) might allow the mechanism generating the static MAE to play a significant role. This possibility is in line with the nearly perfect transfer found by Raymond (1993a), who examined the effect of adaptation on motion coherence thresholds using dynamic random-dot patterns that had less slowly drifting components than those used by Steiner et al. (As discussed later, however, the phenomenon investigated by Raymond and her colleagues appears to have a number of properties distinct from those found in traditional MAE paradigms.)

In addition to the spatiotemporal properties of test stimuli and the methods used to measure MAEs, retinal eccentricity exerts a significant influence on the interocular transfer of the dynamic MAE. Ashida et al. (1997) found that transfer of the MAE (based on duration measured using

counterphasing gratings) was nearly perfect in the central visual field, but was drastically reduced in the peripheral visual field. They suggested that this reduction might partially reflect difficulties in feature tracking or the loss of involuntary attention to the adaptation motion presented in the periphery. This hypothesis is in line with their preliminary results demonstrating that, even in the central visual field, interocular transfer of the dynamic MAE is significantly reduced when subjects cannot attend to the adaptation motion. The role of attention in the dynamic MAE is further addressed in later sections.

5.1.6 Storage

Verstraten et al. (1996) reported almost complete storage of the static MAE when the static test was preceded by a dynamic test. In their experiment, a dynamic random-dot test stimulus was presented immediately after adaptation, for as long as the MAE was perceived, and then the test was replaced by a static pattern. Although the dynamic MAE had disappeared by the time the static pattern was presented, subjects perceived an MAE in the static test which had nearly the same duration as when the static pattern was presented immediately after adaptation. On the other hand, the dynamic MAE was not stored when the dynamic test was preceded by a static test. This is a good demonstration of a dissociation in the time course of recovery from adaptation between static and dynamic MAEs. It has since been replicated using narrow-band stimuli (static and counterphasing gratings; Culham, unpublished observations).

5.1.7 Surround Motion

It is known that the magnitude of the static MAE depends upon the background stimulus presented during adaptation. For instance, the static MAE is enhanced when the moving adaptation stimulus is surrounded by a static pattern (Day and Strelow, 1971), and the static MAE is far stronger when the adaptation stimulus is surrounded by oppositely moving stimuli than when it is surrounded by stimuli moving in the same direction (Murakami and Shimojo, 1995). It has been suggested that relative movement is an important requirement for generating static MAEs (Anstis and Reinhardt-Rutland, 1976; Swanston and Wade, 1992). Nishida and Ashida (1997) recently found that the surround modulation of the dynamic MAE duration was as strong as that of the static MAE duration when measured in the peripheral visual field, but the modulation was reduced (and nearly disappeared for some observers) in the central visual field. This suggests that the role of relative motion is smaller in the generation of dynamic MAEs.

5.1.8 *Separate Mechanisms?*

The discrepancies between the properties of the static MAE and the dynamic MAE are summarized in table 5.1 (along with a number of other factors to be discussed later). These differences are difficult to reconcile with the notion that a single common mechanism is responsible for both types of MAE. Although the only physical difference is the temporal structure of the test stimulus, these discrepancies support the idea that the two types of MAE reveal adaptation at separate sites of visual motion processing. The properties of the static MAE indicate that it reflects adaptation of a mechanism that is spatial frequency-selective, temporal frequency-tuned, and partially monocular. On the other hand, the dynamic MAE appears to reflect adaptation of a mechanism that is non-spatial frequency-selective, velocity-tuned, and completely binocular.

It should be noted, however, that the distinction between the static MAE and the dynamic MAE may not be strict. As mentioned above, even when the test stimulus is a counterphasing grating, the properties of the MAE are similar to those obtained with a static test when the temporal frequency of the dynamic test is lower than 1 Hz (Bex et al., 1996; Ashida et al., 1997). This change in properties with test temporal frequency seems to be accompanied by a change in the appearance of the MAE. When the test stimulus is static or has a low temporal frequency, the MAE is manifest as an apparent gradual shift in the spatial positions of the bright and dark bars. On the other hand, the MAE elicited in a rapidly counterphasing grating appears to have a drift speed almost comparable to that of the individual drifting components composing the test stimulus.

Another notable issue with respect to the multiplicity of MAE mechanisms concerns the effects of retinal eccentricity on the dynamic MAE. The dynamic MAE measured in the periphery is different from that measured in the central visual field, at least in terms of the degree of interocular transfer and the modulatory effect of motion in the surround (Ashida et al., 1997; Nishida and Ashida, 1997). This raises the possibility that two mechanisms are responsible for the generation of the dynamic MAE: one predominates in the central visual field, and the other predominates in the periphery. That the dynamic MAE in the periphery is similar to the static MAE with regard to interocular transfer and the effect of surround motion may imply that the static MAE and dynamic MAE are both mediated by a common mechanism in the peripheral visual field, but by separate mechanisms in the central visual field. The low-level mechanism may generate dynamic MAEs also in the central visual field, but it presumably decays more rapidly than that generated by the high-level mechanism. This could be a reason why the interocular transfer is not perfect when the MAE magnitude is measured in terms of the nulling motion strength.

Table 5.1

Comparison of four types of motion aftereffect (MAE)

	Static MAE	Dynamic MAE ^a	DS threshold elevation	Speed aftereffect
Test contrast	Supra-threshold	Suprathreshold	Threshold	Supra-threshold
Test temporal structure	Static	Dynamic (ambiguous)	Dynamic (drifting)	Dynamic (drifting)
SF selectivity	Yes ¹	No ²	Yes ³	No ⁴
TF/speed	TF ⁵	Speed ⁶	TF and speed ⁷	Speed ⁸
Spatial specificity	Yes ⁹	No ¹⁰		
IOT	Partial ¹¹	Complete ¹² [partial] ¹³	Partial ¹⁴	Partial ¹⁵
Storage effect	Yes ¹⁶	No ¹⁷		
Relative motion	Effective ¹⁸	Ineffective [effective] ¹⁹		
SO motion	Ineffective ²⁰	Effective ²¹	Effective ²²	Effective ²³
FO-SO cross-adaptation	No ²⁴	Yes ²⁵	No ²⁶	Yes ²⁷
Long-range motion	No ²⁸	Yes ²⁹		
Modulation by attention	No ³⁰ Yes ³¹	Yes ³²		
Generated by attentive tracking	No ³³	Yes ³³		

^aDS, direction-selective; SF, spatial frequency; TF, temporal frequency; IOT, interocular transfer; FO, first-order; SO, second-order; Properties of dynamic MAE in brackets are those obtained mainly in the peripheral visual field.

¹Over et al. (1973), Cameron et al. (1992a); ²Ashida and Osaka (1994), Mareschal et al. (1997); ³Pantle et al. (1978), Nishida et al. (1997a); ⁴Thompson (1981); ⁵Pantle (1974), Wright and Johnston (1985); ⁶Ashida and Osaka (1995b); ⁷Sekuler (1975); ⁸Thompson (1981); ⁹Anstis and Gregory (1965); ¹⁰von Grünau & Dubé (1992); ¹¹Moulden (1980); ¹²Nishida et al. (1994); ¹³Ashida et al. (1997); ¹⁴A. T. Smith (1983); ¹⁵A. T. Smith and Hammond (1985); ¹⁶P. Thompson and Wright (1994); ¹⁷Verstraten et al. (1996); ¹⁸Anstis and Reinhardt-Rutland (1976), Swanston and Wade (1992); ¹⁹Nishida & Ashida (1997); ²⁰Anstis (1980), Derrington and Badcock (1985); ²¹McCarthy (1993), Ledgeway and Smith (1994a); ²²Turano (1991), Nishida et al. (1997a); ²³Ledgeway and Smith (1997); ²⁴Nishida and Sato (1995); ²⁵Ledgeway (1994), Ledgeway and Smith (1994a), Nishida and Sato (1995); ²⁶Nishida et al. (1997a); ²⁷Ledgeway and Smith (1997); ²⁸Anstis (1980); ²⁹von Grünau (1986); ³⁰Wohlgemuth (1911); ³¹Chaudhuri (1990); ³²Shulman (1993), Lankheet and Verstraten (1995); ³³Culham and Cavanagh (1994).

5.1.9 Relation to Other Motion Aftereffects

A promising explanation of dynamic MAEs, especially those observed with a counterphasing grating, is that adaptation causes a direction-selective reduction in contrast sensitivity. Since a counterphasing sinusoidal grating can be decomposed into two oppositely drifting components of half-amplitude, a decrease in contrast sensitivity to one component would enhance the relative salience of the other component, possibly resulting in a percept akin to the dynamic MAE. Another phenomenon that indicates a direction-selective sensitivity reduction is the direction-selective elevation of contrast threshold after adaptation to moving stimuli (Sekuler and Ganz, 1963; Pantle and Sekuler, 1969). Levinson and Sekuler (1975) suggest that this threshold elevation and the dynamic MAE, one prominent near-threshold and the other prominent at suprathreshold contrast levels, result from adaptation at the same site. However, unlike the dynamic MAE, threshold elevation is spatial frequency-selective (Pantle et al., 1978; Nishida, Ledgeway, et al., 1997c). Threshold elevation clearly occurs only when the adaptation and test stimuli are both first-order or second-order motion patterns (Nishida et al., 1997c), while the dynamic MAE occurs for all combinations of stimulus type (Ledgeway, 1994). These discrepancies suggest that these two phenomena have separate origins. In addition, the differential efficacy with which second-order motion stimuli give rise to direction-selective threshold elevation and the static MAE suggests that these two phenomena are distinct. It appears that the three types of MAE, the static MAE, the dynamic MAE, and direction-selective threshold elevation, reflect processes occurring at different sites of adaptation.

Another effect of motion adaptation is a change in the perceived speed of moving suprathreshold stimuli (Carlson, 1962). The magnitude of the speed aftereffect depends on the speed of the adaptation stimulus regardless of its spatial frequency (P. Thompson, 1981). In addition, it occurs with second-order patterns even when first-order patterns are used for adaptation or test stimuli (Ledgeway and Smith, 1997). These properties are similar to those obtained with the dynamic MAE (see table 5.1).

The existence and characteristic properties of these different types of MAE could, in principle, be used to elucidate the underlying functional architecture of motion processing in human vision. In the next section, we focus on one area where this approach has recently been widely adopted, namely, the study of second-order motion.

5.2 Motion Aftereffects from Second-Order Motion

Moving objects in the visual world may give rise to drifting contours in the retinal image that are defined in terms of local variations in intensity

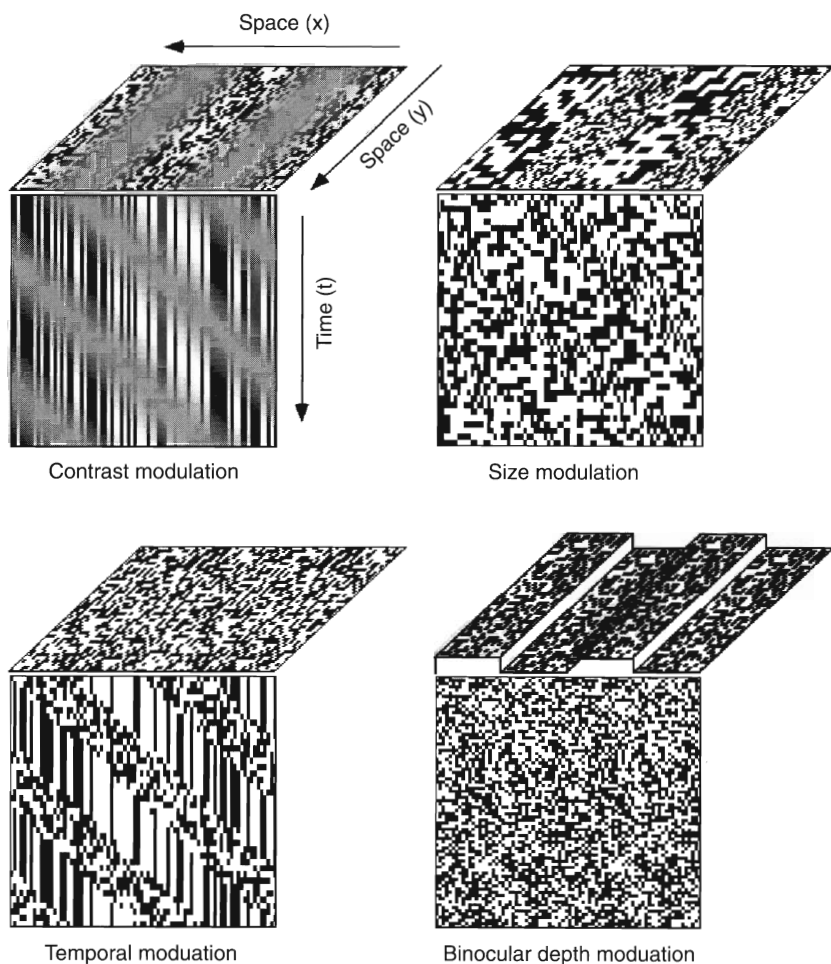
or wavelength (first-order motion) or by variations in more complex derived characteristics (second-order or "non-Fourier" motion) such as local contrast, orientation, and binocular disparity (Chubb and Sperling, 1988; Cavanagh and Mather, 1989). The study of second-order motion is important theoretically because intensity-based motion analyzers, which have been modeled as quasi-linear filters that are oriented in space-time or spatiotemporal frequency (e.g., Adelson and Bergen, 1985), should be incapable of detecting second-order motion. Yet stimuli that give rise to second-order motion (figure 5.2), such as a drifting modulation of the contrast of a stationary luminance grating (carrier), typically produce vivid percepts of movement, and this poses a problem for conventional motion-detecting schemes.

5.2.1 Theories and Models of Second-Order Motion Perception

Since second-order motion stimuli undoubtedly contain "features" (e.g., texture edges and boundaries), motion perception could, in principle, be mediated primarily by a high-level, cognitive strategy involving the explicit tracking or matching of image features over space and time (Anstis, 1980; Braddick, 1980; Cavanagh, 1992). However, several computational models have been proposed which detect second-order motion on the basis of low-level, preattentive mechanisms without the need to track, or calculate correspondences between, features in the image. For instance, it has been suggested that first-order and second-order motion are detected by a single (common) low-level mechanism, an approach which is exemplified by the spatiotemporal gradient model of Johnston, McOwan, et al. (1992). An alternative class of model suggests that first-order and second-order motion are detected by independent (distinct) low-level mechanisms in the visual system (e.g., Derrington and Badcock, 1985, 1986; Chubb and Sperling, 1988, 1989a,b; Badcock and Derrington, 1989; Wilson, Ferrera, et al., 1992; Werkhoven, Sperling, et al., 1993; Zanker, 1993; Fleet and Langley, 1994; Nishida et al., 1994; Nishida and Sato, 1995). These models typically exploit the principle that a nonlinear transformation (e.g., rectification) applied to the luminance profile of many second-order motion stimuli is sufficient to expose their motion to analysis by sensors which utilize the same principles as those used for detecting first-order motion. Adaptation techniques have proved to be useful tools for discriminating between these different classes of model.

5.2.2 Motion Aftereffects Measured with Static, Second-Order Test Patterns

It is well established that following adaptation to stimuli that give rise to second-order motion, the presentation of a stationary test pattern typically fails to elicit a compelling MAE. Indeed, adaptation to several

**Figure 5.2**

Space (x)–space (y) and space (x)–time (t) plots of typical second-order motion stimuli that have been used to investigate motion aftereffects. In each case the stimulus can be constructed by modulating either the contrast (top left), size (top right), flicker rate (bottom left), or binocular disparity (bottom right) of the random elements constituting a two-dimensional noise field (carrier) with a drifting, one-dimensional, sinusoidal, or square-wave profile. With the exception of the contrast-modulated image, the carrier is dynamic rather than static (although it is possible to employ a dynamic carrier in a contrast-modulated pattern). The depicted stimuli are members of a class of motion patterns termed “drift-balanced” by Chubb and Sperling (1988) that, in principle, do not contain any net first-order motion energy.

varieties of second-order motion stimuli, including kinetic edges (Anstis, 1980), drifting beat patterns (Derrington and Badcock, 1985), contrast-modulated sinusoidal gratings (Turano and Pantle, 1985; McCarthy, 1993), and contrast-modulated noise (Ledgeway and Smith, 1993), does not produce a discernible MAE when tested with static second-order patterns. However, some second-order motion stimuli such as disparity-defined (cyclopean) gratings (Patterson et al., 1994; Bowd et al., 1996b; see also chapter 4) and contrast-reversing bars (Mather, 1991) do generate measurable MAEs with stationary test patterns, although their magnitudes (quantified in terms of duration) are considerably weaker than those resulting from adaptation to comparable first-order (luminance-defined) motion.

Explanations as to why second-order motion is so ineffective in generating MAEs with stationary test stimuli have tended to focus on either the properties of motion detectors in the visual system or the physical characteristics of the motion stimuli themselves. In terms of the former, it is possible that second-order motion is encoded principally by feature-tracking or feature-matching processes, since adaptation to motion stimuli which are believed to favor such processes does not generate compelling MAEs with stationary patterns (e.g., Banks and Kane, 1972; Anstis and Mather, 1985). Although feature-based strategies may be used, at least some of the time, for the detection of second-order motion (A. T. Smith, 1994a; Lu and Sperling, 1995b), second-order motion patterns support several phenomena which are widely held to preclude explicit feature-encoding strategies (see Cavanagh and Mather, 1989, for a review).

For instance, observers can readily identify the direction of displacement of dense, random-dot patterns in which the dots are defined by contrast variations (e.g., Nishida, 1993) and this suggests that second-order motion, like first-order motion, is likely to be detected most of the time using low-level mechanisms. If second-order motion perception is mediated primarily by low-level mechanisms, as suggested by several current models, then the failure to find compelling MAEs with stationary test stimuli implies that these mechanisms are distinct, and may have different temporal properties, from those that detect first-order motion (Derrington and Badcock, 1985; Badcock and Derrington, 1989; Pantle and Turano, 1992; Derrington, Badcock, et al., 1993; McCarthy, 1993).

However, this proposition is complicated by the finding that when first-order motion stimuli such as luminance gratings are equated for the visual system's relatively poor sensitivity to second-order motion (A. T. Smith, Hess, et al., 1994) or the presence of a salient carrier, MAEs measured with stationary test stimuli are either degraded or absent (A. T. Smith, Musselwhite, et al., 1984; Pantle, Pinkus, et al., 1991; Ledgeway and Smith, unpublished observations cited in A. T. Smith, 1994b; McCarthy,

1993). A stationary carrier may serve to reinforce the immobility of the test stimulus and mask or override weak motion signals derived from differential activation of adapted and unadapted motion detectors (Ledgeway, 1994). In addition, carriers give rise to substantial first-order motion in all second-order motion stimuli (although this may be incoherent or drift in a direction different from that of the second-order information) which may act as a source of external noise during adaptation and degrade signal fidelity in the pathways that encode the direction of second-order motion. These possibilities could account for the finding that cyclopean gratings evoke demonstrable MAEs only when the stationary second-order spatial structure of the test pattern is defined by disparity modulation of a dynamic random-dot carrier, and a relatively long adaptation duration (e.g., 2 minutes) is employed (Patterson et al., 1994; Bowd et al., 1996b). However, an alternative explanation of this phenomenon will be discussed later in the context of MAEs measured with dynamic test patterns.

5.2.3 Motion Aftereffects Measured with Dynamic, Second-Order Test Patterns

In contrast to the weak or nonexistent MAEs measured with stationary test patterns, it has been found that adaptation to second-order motion results in substantial aftereffects when dynamic (flickering or drifting) test stimuli are employed. For example, following adaptation to a drifting sinusoidal modulation of the contrast of a random noise field or luminance grating, a directionally ambiguous, second-order test stimulus appears to drift coherently in the direction opposite to adaptation (McCarthy, 1993; Ledgeway and Smith, 1994a). Using a nulling method Ledgeway (1994) found that this MAE could be abolished by manipulating the relative amplitudes of the two oppositely drifting sinusoidal components present in a counterphasing test stimulus such that the amplitude of the component drifting in the same direction as adaptation was approximately twice that of the other component. An MAE of comparable magnitude was found when the adaptation and test stimuli were both first-order, luminance-modulated noise patterns presented at the same multiple of direction threshold as the contrast-modulated stimuli. MAEs measured with dynamic second-order stimuli are not limited to periodic, contrast-defined patterns because adaptation to the stereoscopic motion of disparity-defined random-dot patterns has been found to bias the perceived drift of similar second-order patterns by as much as 20 degrees away from the adaptation direction (Patterson and Becker, 1996). Other manifestations of the MAE have also been reported with second-order motion stimuli. For instance, prolonged exposure to the motion of contrast-modulated noise produces direction-selective increases and decreases in the perceived

speeds of drifting test patterns (Ledgeway and Smith, 1997), similar to those reported for luminance gratings (P. Thompson, 1981). Second-order MAEs are not confined to suprathreshold stimuli, since postadaptation detection and motion thresholds for contrast-modulated gratings (Turano, 1991), noise (Ledgeway and Smith, 1992), and beat patterns (Holliday and Anderson, 1994) are elevated relative to those measured in the absence of prior adaptation to motion. These aftereffects exhibit direction selectivity in that thresholds are elevated more for stimuli that drift in the same direction as adaptation than for stimuli that drift in the opposite direction.

An important property of MAEs measured with dynamic test stimuli is that typically they are insensitive to the particular variety of motion presented during adaptation. That is, adaptation to first-order motion affects the perceived direction and speed of second-order motion, and vice versa (e.g., Patterson et al., 1994; Nishida and Sato, 1995; Patterson and Becker, 1996; Ledgeway and Smith, 1997). Indeed, several studies (Turano, 1991; Ledgeway, 1994; Ledgeway and Smith, 1994a) have reported that "cross-adaptation" between first-order and second-order motion patterns produces robust aftereffects of similar magnitude to those found when the adaptation and test stimuli are either both first-order or both second-order motion patterns, particularly when the two varieties of motion are equated for visibility. These phenomena suggest that dynamic test stimuli may provide a more sensitive technique for measuring postadaptation changes in sensitivity to second-order motion than stationary test patterns. Nonetheless, as discussed previously, there is evidence (e.g., Nishida et al., 1994) that aftereffects measured with static and dynamic test stimuli may have qualitatively different bases. For example, Nishida and Sato (1992, 1995) found that following adaptation to a beat pattern in which first-order (luminance-defined) and second-order (contrast-defined) motion drifted simultaneously in opposite directions, a static beat or first-order sinusoidal grating appeared to drift in the direction opposite to the first-order adapting motion. However, the direction of the MAE reversed when the test stimulus was a first-order grating counterphasing at 2 Hz. Thus, it is possible that static and dynamic test stimuli probe the properties of different mechanisms in the pathways that process motion information.

Given the existence of cross-adaptation effects between first-order and second-order motion, great care needs to be taken when investigating conventional MAEs with test patterns that contain static second-order spatial structure in conjunction with a dynamic carrier, such as cyclopean gratings (e.g., Patterson et al., 1994). This is because following adaptation to second-order motion it is possible that any perceived opposite drift of the test stimulus is due entirely, or at least in part, to an aftereffect induced in the dynamic, first-order carrier as a result of cross-adaptation.

Evidence for this possibility comes from the finding (Ledgeway, unpublished observations) that adaptation to a drifting stimulus composed of either contrast-modulated or luminance-modulated dynamic noise produces robust MAEs of comparable magnitude when an unmodulated dynamic noise pattern (carrier) serves as the test stimulus. Specifically, the first-order and second-order adaptation patterns were both presented at approximately 4.5 times direction-identification threshold and the modulation signal in each case was a vertically oriented, 1 cycle/degree sinusoid drifting at 5 Hz. The carrier was spatially two-dimensional and replaced with a new stochastic sample at a rate of 32 Hz to create dynamic noise. Following 2 minutes of adaptation to motion, the modulation signal was removed and the resulting unmodulated dynamic noise field appeared to drift coherently in the direction opposite to adaptation for approximately 5 to 6 seconds, irrespective of whether the adaptation stimulus was a first-order or second-order motion pattern. Thus, attempts to isolate after-effects exhibited by the second-order structure of a test stimulus in which the carrier is dynamic need to ensure that measures of performance are not contaminated by the presence of MAEs in the first-order carrier itself.

5.2.4 Interpreting Cross-Adaptation Effects Within Models of Motion Processing

The extent to which MAEs measured with dynamic first-order and second-order test stimuli reflect the relative contributions of high-level and low-level motion processes is at present largely unresolved. However, as we will see later in this chapter, in the context of first-order motion there is evidence that high-level, attention-based, feature-tracking processes may mediate, or at least exert a strong modulatory influence upon, some manifestations of the MAE observed with dynamic test patterns (e.g., Culham and Cavanagh, 1994) and it is possible that similar processes also operate within the domain of second-order motion. Nevertheless, it seems doubtful that MAEs measured with dynamic stimuli can be attributed exclusively to the operation of high-level motion-detecting processes, because under a condition believed to selectively favor feature-based motion detection (i.e., the introduction of an interstimulus-interval [ISI] of 60 to 90 msc between successive updates of the spatial position of the adaptation stimulus), the MAE measured with a directionally ambiguous, first-order or second-order test pattern is severely degraded (A. T. Smith and Ledgeway, 1994). This implies that low-level, motion-detecting mechanisms play a direct role in producing MAEs with dynamic first-order and second-order patterns. In terms of current models of motion detection, cross-adaptation effects between first-order and second-order motion appear, at least superficially, strongly to favor models employing a common mechanism for detecting both varieties of motion (e.g., Johnston et

al., 1992). However, it is also possible to accommodate cross-adaptation effects within models incorporating separate detection mechanisms for first-order and second-order motion. If the outputs of distinct first-order and second-order motion detectors are subsequently pooled to compute the resultant direction (and speed) of image motion (e.g., Wilson et al. 1992), then adaptation of motion-sensitive mechanisms either prior to or at the stage at which motion signals are integrated will bias the perceived direction (and speed) of subsequently presented dynamic test patterns.

Although cross-adaptation effects considered in isolation do not allow differentiation between current motion models, some adaptation phenomena may be difficult to reconcile with models that utilize a common mechanism to encode both first-order and second-order motion. For example, Hammett, Ledgeway, et al. (1993) found that observers consistently reported transparency (simultaneous movement in opposite directions) when presented with a beat pattern in which first-order (luminance-defined) and second-order (contrast-defined) motion drifted in opposite directions. Following adaptation to a first-order sinusoidal grating that periodically reversed its drift direction, the percept of transparency was typically abolished and the beat pattern appeared to drift in the direction of the second-order motion. These results clearly demonstrate that adaptation to first-order motion may, under some circumstances, selectively desensitize mechanisms responsive to first-order motion but not second-order motion.

Similarly, Nishida and Sato's (1995) finding that adaptation to a beat pattern, analogous to that used by Hammett et al. (1993), can produce MAEs in opposite directions is commensurate with simultaneous adaptation of separate first-order and second-order motion-detecting mechanisms. A recent study by Nishida et al. (1997c) also has a bearing on this issue. These authors directly compared the spatial-frequency selectivity of first-order and second-order MAEs by measuring direction thresholds for sinusoidal luminance gratings and contrast-modulated noise both prior to and following adaptation to motion. Adaptation to either first-order or second-order motion, of comparable visibility, elevated thresholds for detecting the same variety of motion in a manner that was both direction-selective and spatial frequency-selective. That is, band-limited losses in sensitivity were found that were maximal when the adaptation and test stimuli had the same drift direction and spatial frequency. Importantly, although cross-adaptation effects were sometimes observed, these failed to exhibit spatial-frequency selectivity. These results clearly support the notion that, at least initially, first-order and second-order motion are detected by separate but qualitatively similar mechanisms in the visual system that are each selectively sensitive to a limited range of spatial frequencies. Moreover, convergent evidence from a diverse range of

psychophysical (e.g., Doshier, Landy, et al., 1989; Landy, Doshier, et al., 1991; L. R. Harris and Smith, 1992; Mather and West, 1993; Ledgeway and Smith, 1994b; Edwards and Badcock, 1995; Nishida et al., 1997b; A. T. Smith and Ledgeway, 1997), electrophysiologic (Zhou and Baker, 1993, 1994, 1996), and neuropsychological (Vaina, Le May, et al., 1993) studies is also consistent with the existence of separate mechanisms for encoding each variety of motion. Thus, given the current balance of evidence, cross-adaptation effects are likely to reflect integration of first-order and second-order motion signals at relatively late stages of motion processing, as suggested by some motion models (Wilson et al., 1992; Zhou and Baker, 1993; Nishida and Sato, 1995).

In summary, adaptation studies have provided valuable insights into the functional architecture of the mechanisms that encode first-order and second-order motion, and many adaptation phenomena can be readily interpreted using current theoretical models of motion processing (a topic that is addressed in detail in chapter 7). However, there remain many issues concerning adaptation and second-order motion that warrant further investigation. For instance, little is known about the precise mechanisms by which carriers influence the quality and magnitude of after-effects, or about the relationships and possible interactions between high-level, cognitive strategies and low-level processes in the generation and control of MAEs evoked with dynamic, second-order test patterns.

5.3 Attentional Influences on Motion Aftereffects

Since Wohlge-muth (1911), the MAE has traditionally been thought to be a "preattentive" phenomenon; that is, it has been presumed to result from early and automatic motion processes which do not depend on the observer's state of attention. However, growing evidence is emerging to challenge this view. First, a number of investigators have discovered that the perceived duration, direction, and strength of MAEs may indeed be affected by the degree of attention allocated to the adapting stimulus. Second, newer dynamic tests have revealed MAEs for higher-order stimuli which may be processed through an attention-based motion process. Here we review the evidence for attentional modulation of MAEs and consider the possibility that attentional shifts themselves may lead to MAEs which are qualitatively different from the classic MAE. Further, we consider possible neural substrates for the MAE and its attentional influences.

5.3.1 Attentional Modulation of Motion Aftereffects

Initial investigations by Wohlge-muth (1911) suggested that a subject's attentional state had no effect on the MAE. He found no diminution of

the static MAE when subjects listened to reading, added numbers read aloud, or rapidly read or added numbers presented in the center of the adapting display. More recently, however, Chaudhuri (1990) reported marked reductions in the duration of the MAE using a similar paradigm. Chaudhuri found that MAE duration was reduced by 70 percent when observers performed a foveal attention-distracting task (monitoring a letter stream for the occurrence of digits) on a moving textured background. The foveal attention task had only minor effects on subjects' eye movements, which did not appear to account for the drastic reduction in the MAE. No shortening of MAE duration was found when the task required attention to a diffuse area encompassing the moving stimulus, suggesting that attention must be allocated to a different location for a reduction to occur. Although it is not clear why these two sets of experiments differ, a number of other preliminary reports have corroborated Chaudhuri's report of attentional effects on MAEs (Zhou and Chen, 1994; Boutet, Rivest, et al., 1996; Georgiades and Harris, 1996).

A number of recent experiments have also shown that attention to one of two moving components can bias subsequent judgments of motion. Shulman (1993) replicated Chaudhuri's attentional distraction result using a more powerful opponent technique that involved comparing the perceived direction of the MAE following attention to one direction of motion vs. the opposite direction. In his display, two sets of four grating patches were arranged concentrically and rotated in opposite directions while subjects attended to one set vs. the other. Afterward, an ambiguously rotating test set positioned midway between the two adaptation sets was more likely to be seen rotating opposite to the attended direction. Note that like flicker MAE, the test was directionally ambiguous, suggesting that dynamic MAEs are also modulated by attention.

Attention can affect the MAE not only when it is allocated to different regions of space but also when it is allocated selectively within a single region. Lankheet and Verstraten (1995) presented subjects with two superimposed and transparent random-dot patterns moving in opposite directions and asked subjects to attend to one of the two directions. Following attention to one direction, subjects were more likely to see a dynamic noise pattern moving in the opposite direction, an effect that was quite strong when quantified using a signal/noise nulling procedure. Further, Boutet et al. (1996) found that static MAE strength was enhanced when subjects made attentionally demanding judgments of the motion of the adaptation pattern relative to when they made similar judgments about its luminance or color.

One further study suggests that attention may also modulate MAEs for rotation in depth. Shulman (1991) used his opponent paradigm to demonstrate that attention to one of two squares rotating in depth biased the

percept of a square in parallel projection which ambiguously rotated in depth.

Thus it appears that attention can enhance the processing of motion for an attended region of motion or even for the attended direction of motion within a region. In addition to this attentional modulation of MAEs, a second intriguing possibility is the attentional mediation of MAEs.

5.3.2 Attentional Mediation of Motion Aftereffects

In addition to attentional enhancement of MAEs for attended regions, attributes, or directions, recent reports suggest that MAEs may also result from prolonged attentional shifts in the same direction. That is, attention to visual features may also produce a reversal in perceived direction of ambiguous stimuli. One key question is whether these attentional after-effects result from attentional shifts per se, or simply from effects similar to the attentional modulations described above. Originally, stimuli in discontinuous apparent motion were said not to generate MAEs (Anstis and Mather, 1985; Mather, Cavanagh, et al., 1985) and this was considered one of the criteria for distinguishing short- and long-range motion systems (Anstis, 1980). However, this distinction was questioned when von Grünau (1986) reexamined MAEs from long-range motion using an ambiguous test stimulus, and did indeed find MAEs. In von Grünau's paradigm, subjects adapted to the discontinuous apparent motion of a large grating patch that was repeatedly displaced from one side of the display to another.

Subsequently, observers viewed a counterphasing test grating. Adaptation produced a bias such that observers often reported that the test grating appeared to move opposite to the direction of adaptation, though the effect was somewhat weaker for adaptation to discontinuous motion than to smooth drifting motion. This flicker MAE from discontinuous motion survived dichoptic presentation of the stimuli, indicating that short-range motion mechanisms (known to be disrupted by dichoptic presentation; see Anstis, 1980; Braddick, 1974) were unlikely to be responsible. Furthermore, the MAE occurred even using test locations that did not overlap with the successive positions in which the discontinuous adapting grating was presented—a result that was surprising at the time, because static MAEs are typically location-specific.

Von Grünau's demonstration of MAEs from apparent motion was initially attributed to the greater sensitivity of directionally ambiguous tests, compared to traditional static tests (Cavanagh and Mather, 1989; von Grünau and Dubé, 1992). In retrospect, given that dynamic MAE tests have different properties than static tests and may involve later stages of motion processing, it may instead be that dynamic tests are susceptible to

higher-order factors such as attention that have little or no influence on the static MAE.

Culham and Cavanagh (1994, 1995) have proposed that the flicker MAE, but not the static MAE, can be triggered by shifts in attention that are used to keep track of a moving target's position. This high-level attentive tracking process is said to form the basis of a second motion system distinct from early detector-based motion processes, and to be implicated in the perception of "long-range" or apparent motion (Wertheimer, 1912; Cavanagh and Mather, 1989; Cavanagh, 1991, 1992; Lu and Sperling, 1995a).

Culham and Cavanagh (1994) investigated the effect of attentive tracking on static and flicker MAEs. In order to generate attentive tracking in the absence of net low-level motion signals, they presented a radial grating in counterphase (equal motion energy in opposite directions). Observers used attention to track the bars of the grating in one direction, clockwise or counterclockwise, while fixating on a central bull's-eye. Thus the stimulus was the same for both conditions; only the subjects' instructions differed. Subsequently, static and flicker MAEs were measured using a nulling technique. Attentive tracking of the counterphase grating produced a flicker MAE but not a static MAE. Furthermore, flicker MAEs following attentive tracking were relatively independent of the adapting contrast of the grating, suggesting that any stimulus visible enough to be tracked would produce a flicker MAE. In comparison, passive viewing of smooth motion produced both MAEs which were highly dependent on the adapting contrast.

In a subsequent experiment, Culham and Cavanagh (1995) investigated the retinotopic specificity of MAEs from smooth motion and attentive tracking. Although von Grünau and Dubé (1992) found flicker MAEs even for unadapted test locations, this experiment further investigated whether MAEs were specific to the global direction of rotation or the local direction of motion components. In the display, shown in figure 5.3A, a radial adapting grating was presented in a central ring and the test ring was presented in the periphery to produce a region of overlap. This configuration yielded two possible outcomes for the direction of MAEs: first, "local MAEs" to adaptation could result if linear motion detectors within the region of overlap led to a locally opposite MAE (e.g., for counterclockwise adaptation, local downward motion in the overlap would produce an upward MAE, biasing the observer to see rotation in the same direction as adaptation, counterclockwise); second, "global MAEs" could result from the adaptation of rotation-selective mechanisms (e.g., counterclockwise adaptation would produce a clockwise MAE). As shown in figure 5.3B, attentive tracking produced flicker MAEs opposite

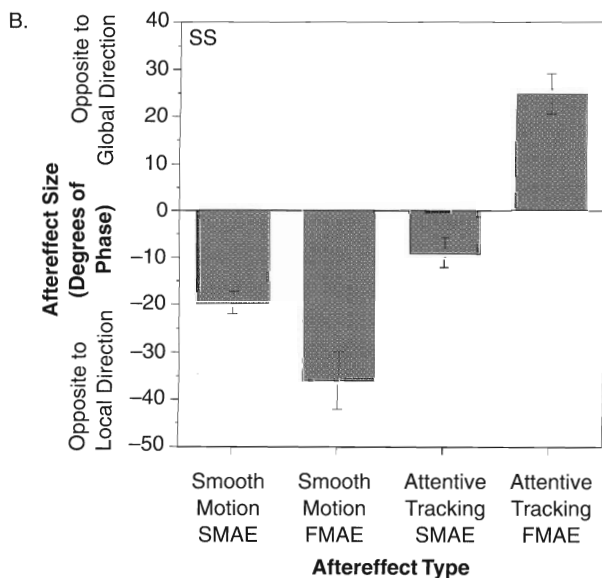
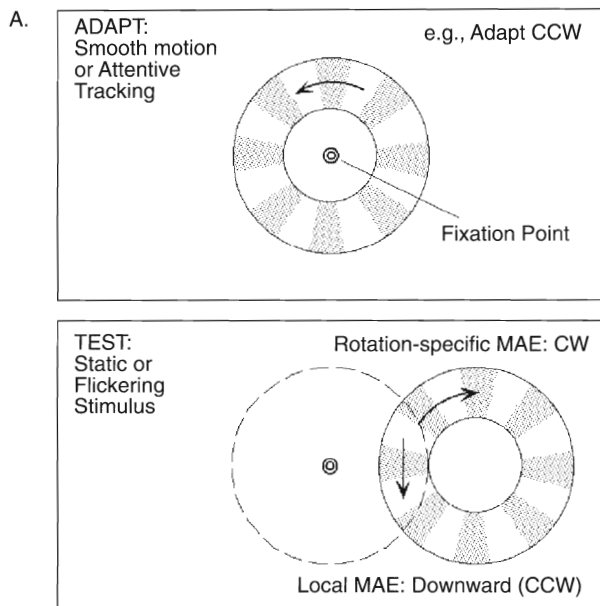


Figure 5.3

Stimulus used to examine retinotopic specificity of attentional aftereffects. Subjects either passively viewed smooth motion or attentively tracked a counterphase grating. The test stimulus was displaced from the adapt to produce a region of overlap which yields two

to the global rotation of the tracked bars. In comparison, attentive tracking led to a negligible static MAE in the local direction. Adapting to smooth motion produced large local effects for both static and flicker MAE tests. This interaction suggests that attentive tracking acts at a relatively high level of motion perception in which the analysis can be rotation-specific, independent of position, and selective to tracked targets as opposed to physical motion (which was balanced in the counterphasing adapting stimulus).

One possibility raised by these results is that attentive tracking of visible features, rather than low-level motion energy, may be responsible for flicker MAEs observed with a variety of stimuli, including apparent motion (von Grünau, 1986) and possibly second-order motion (Nishida and Sato, 1995). For example, the results of Nishida and Sato, showing that second-order motion patterns (e.g., defined by texture or depth) generate flicker MAEs, may indicate that these attributes are processed by higher-order detectors which affect only late stages of motion processing tapped by the flicker test. Alternatively, subjects may be attentively tracking the form-defined bars of the adapting patterns, in which case attention alone could produce a flicker MAE. This is not to say these results are necessarily attention-based, though further evidence is required to distinguish the two possible accounts. It may also be that not all second-order patterns are processed the same way, with some types, such as texture patterns, processed by detector-based systems, and other types, such as depth-defined or motion-defined patterns, processed by attention-based systems (Cavanagh, 1995).

Whether or not attention is primarily responsible, it appears to have a strong influence on flicker MAE results using higher-order stimuli. For example, two groups have examined flicker MAEs following adaptation to missing fundamental patterns (Bex, Brady, et al., 1995) or to second-order beat patterns (Ashida and Osaka, 1995a), both of which include two opposing motion components—low-level motion energy in one direction and feature displacement in the other. Bex et al. reported that when no ISI is present, the flicker MAE appears opposite to motion energy; but when an ISI is used, the MAE appears opposite to feature motion. Ashida and Osaka (1995a) reported similar findings, but noted that when an ISI is

Figure 5.3 (continued)

potential outcomes for motion aftereffects (MAEs): local MAEs opposite to the linear motion within the region of overlap, or global MAEs opposite to the overall direction of rotation. The phase shifts necessary to null the MAEs are shown for one observer in panel B. Positive values indicate outcomes consistent with the global MAE; negative values indicate a local MAE. Vertical bars represent ± 1 SE. CCW, counterclockwise; CW, clockwise; SMAE, Static MAE; FMAE, Flicker MAE.

used the MAE was observed only if subjects had to make a judgment about the perceived direction of the adaptation pattern, suggesting that attention was necessary to produce the high-level effect.

5.3.3 Neural Correlates of Attentional Effects on Motion Aftereffects

Attentional influences on the MAE may provide clues as to the neural basis of interactions between attention and motion processing. One very likely candidate for attentional modulation of MAEs in humans is area MT (middle temporal visual area). Specifically, Tootell and his colleagues (1995) have shown that the time course of the functional magnetic resonance imaging (fMRI) signal in the region around "greater MT" (possibly including the medial superior temporal area, or MST) is closely correlated with the subjective strength of the MAE as it decays (see chapter 6). While other motion areas also produced an enhanced magnetic resonance (MR) signal during the MAE, the time course of the response did not match that of the MAE percept (Dale, Reppas, et al., 1996). Furthermore, two other neuroimaging studies have demonstrated stronger activation in greater MT when subjects made judgments of motion than when they made judgments of form or color (Corbetta, Miezen, et al., 1990, 1991), or when they attended to moving dots rather than superimposed stationary dots (O'Craven, Rosen, et al., 1997). A number of physiologic studies have suggested that MST rather than MT is susceptible to "top-down" influences (Wurtz, Richmond, et al., 1984; Newsome, Wurtz, et al., 1988; Ferrera, Rudolph, et al., 1994). However, one elegant study by Treue and Maunsell (1996) demonstrated attentional influences in both areas depending on the nature of the task. Both regions had cells which showed modulation when monkeys attended to a moving dot within vs. outside their receptive fields. One additional possibility is that the locus of attentional modulation depends on the particular motion task. Based on fMRI experiments, Watanabe and Miyauchi (in press) have proposed that area V1 is modulated by attention to motion components, while MT is modulated by attention to integrated motion.

As for the locus of attentional mediation of MAEs, a number of interesting possible neural substrates exist. Given the evidence that attentive tracking leads to global rather than local effects (Culham and Cavanagh, 1995), one possibility is that attentive tracking acts at a level such as MST, which has cells selective to rotation direction and is susceptible to extraretinal influences. However, an intriguing alternative is that attentive tracking and the phenomena that it governs, such as apparent motion and its aftereffects, reflect the activity of higher motion-processing stages. Here, the most likely candidate is the posterior parietal lobe, which has cells responsive to the inferred motion of occluded stimuli (Assad and Maunsell, 1995), has many of the properties necessary for apparent

motion correspondence (see Dawson, 1991, for a review), and has been demonstrated to be active while viewing attentive tracking displays (Culham, Cavanagh, et al., 1997). Although further research is required to investigate such a suggestion, this case illustrates that not only may attentional factors be used to understand MAEs but, conversely, MAEs may be used to investigate other higher-order processing such as visual attention.

5.4 Motion Aftereffects from Multivectorial Motion

In previous sections we have dealt with MAEs from simple stimuli, mostly sinusoidal gratings defined by first- or second-order attributes. However, a brief reflection on everyday instances of visual motion suggests that we commonly experience complex interactions between motions. For example, real-world stimuli contain diverse spatiotemporal components, and multiple sets of items can move transparently through each other, sometimes at different depths, as when one looks through falling rain to notice trees bending in a storm. Furthermore, physiologic studies find that neurons in various direction-selective regions show different responses to complex motion stimuli. Thus, restricting our study to simple univectorial motion would neglect the important interactions which can occur between components comprising complex motions. In this section, the study of multivectorial motion is reviewed to investigate not only the neural basis of MAEs but also the levels at which moving components become integrated.

Multivectorial motion refers to the situation in which a given spatial location contains motion in more than one direction. For the most part, investigations of multivectorial motion have been restricted to two motion directions, although in principle more directions could be used. Thus, in discussing multivectorial motion, we will be primarily talking about bivectorial motion. As with other kinds of motion, adaptation to multivectorial motion will produce an MAE. Curiously, however, MAEs induced by adaptation to multivectorial motion are generally not themselves multivectorial. That is, adaptation to multivectorial motion usually generates a univectorial MAE with a direction approximately opposed to the vector sum of the adapting motions. This is the case for bivectorial motion when the two adapting directions are presented in temporal alternation (e.g., Riggs and Day, 1980) or simultaneously (e.g., Mather, 1980; Movshon, Adelson, et al., 1985), and occurs irrespective of whether the adapting motions are perceived to move coherently (as may occur when superimposed gratings form a coherent plaid) or transparently (as occurs when two sheets of random dots are superimposed). Given the unitary nature of MAEs produced by adaptation to this kind of motion,

terming them "multivectorial MAEs" is potentially misleading. In the discussion that follows, the term "MAE of multivectorial motion" is used in preference.

5.4.1 Motion of Aftereffects of Transparent Motion

It is surprising indeed that the MAEs of multivectorial motion are unidirectional even in the case of adaptation to completely transparent motion—even though two clearly different directions are perceived during adaptation, the MAE is opposite to a direction that has never been perceived, while the perceived directions apparently produce no MAEs. What happens to the MAEs presumably generated by the component motions?

This question has recently been investigated in a series of studies by Verstraten and colleagues. They used moving random pixel arrays (RPAs—two-dimensional arrays of dots in which each pixel is randomly assigned to be black or white and the pixels are shifted to generate motion). Two RPAs with directions 90 degrees apart were combined and formed a transparent display in which the individual RPA motions were visible. MAE direction was measured on static RPAs for various speed combinations of the two motions. For unequal speeds, MAE direction could be predicted by the vector sum of the components weighted by a measure of sensitivity to the individual components, such as MAE duration and directional sensitivity. Since MAE direction did not change over time, even when one component would be expected to elicit only a brief MAE, Verstraten and colleagues suggested that the MAE is not determined for each component and then combined, but created only at or after an integration stage, perhaps at MT or later (Verstraten et al., 1994a).

A further study (Verstraten et al., 1994b) found that adaptation to unidirectional motion is stored and used in combination with adaptation to subsequent orthogonal motion. The earlier stimulus had a longer influence on the direction of the combined MAE than would be expected from its MAE when measured in isolation, providing additional evidence that the MAE is formed at or after the site where components become integrated. The region over which the component motions influence the integrated MAE appears to be substantially larger than that of the spatial integration area involved in perceiving coherent motion, again suggesting that MAE generation is at a higher level (van Wezel, Lankheet, et al., 1996). Furthermore, the direction of the MAE to transparent motion also depends on the test stimulus type (static noise vs. dynamic noise of various temporal frequencies). One account of these results (Verstraten, 1994) is that they reflect interactions between spatiotemporal filters that are located before the site at which MAEs are generated. This explanation is consis-

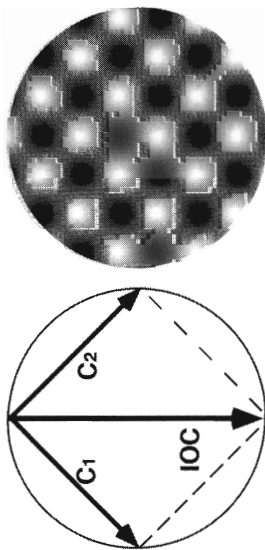
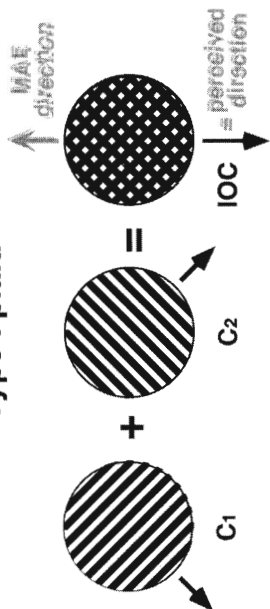
tent with the results of physiologic studies that fail to show fatiguing effects at early visual levels that last long enough to support the durations of MAEs observed psychophysically (Barlow and Hill, 1963; Hammond, Mouat, et al., 1986, 1988; Hammond, Pomfrett, et al., 1989). This account is not incompatible with the idea that MAEs are generated at different levels of processing (e.g., Nishida and Sato, 1995), but states only that the relevant levels may be relatively late in the sequence of processing.

5.4.2 Motion Aftereffects of Coherent Motion

A novel kind of multivectorial stimulus known as a plaid was introduced by Adelson and Movshon in 1982 (but see also Exner, 1887, and Borschke and Heschles, 1902, cited in chapter 1). Plaids are stimuli constructed from superimposed sinusoidal or rectangular grating components drifting in different directions (see figure 5.4). Unlike multivectorial stimuli composed of random-dot patterns, coherent motion is often observed with plaid stimuli (Kim and Wilson, 1993)—the motions of the two component gratings are no longer perceived as independent and a new entity, the plaid pattern, is seen to move rigidly in a third direction. With non-coherent plaids, the component gratings are seen clearly to slide transparently over each other. Movshon et al. (1985) proposed a two-stage model to account for these observations based on their neurophysiologic observations that two kinds of direction-selective cells exist in area MT. The first kind, termed component-selective, responded maximally to the motion of the plaid's components in a manner similar to V1 cells, whereas a minority of cells (approximately 25 percent) were found to be pattern-selective and responded maximally to the coherent motion of the plaid pattern itself. In the first stage of their model, direction-selective units in V1 detect the motion of the component gratings and feed their signals to area MT where, if coherence conditions are met, the component motions are integrated into a coherently moving plaid.

Movshon et al. (1985) proposed that the perceived direction of coherent plaids is determined by the intersection of constraints (IOC) solution, a geometric solution to the inherent directional ambiguity of the plaid's one-dimensional (1-D) components (Adelson and Movshon, 1982; Fennema and Thompson, 1979). Consistent with this model, plaids are usually seen to move in a direction corresponding to that predicted by the IOC (figure 5.4, left side). Plaids in which the component motion vectors both lie to the same side of the resultant vector do not obey this rule (these are so-called type II plaids; figure 5.4, right side; see also Fererra and Wilson, 1987, 1990). Perceived direction in type II plaids is biased toward the component directions and away from the IOC prediction (see figure 5.4, right side). Nonetheless, the MAE is again unitary and is seen to move in the direction opposite to the perceived plaid direction for both

Type I plaid



Type II plaid

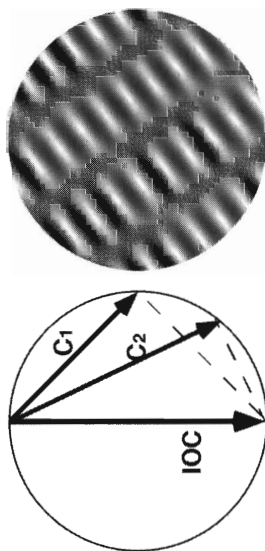
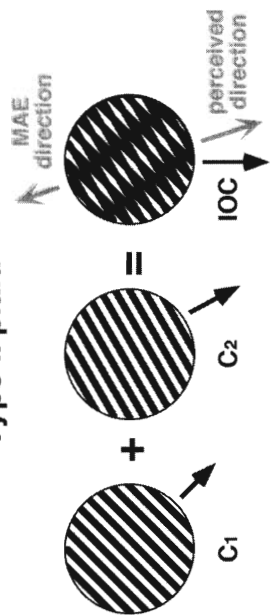


Figure 5.4

Plaid stimuli are formed by the addition of two independently moving gratings (C1 and C2). Following Ferrera and Wilson's (1987) classification, examples are shown of type I plaids (for which component vectors lie on each side of the intersection of constraints [IOC] resultant) and type II plaids (for which both component vectors lie to one side of the IOC resultant). The lower rows show the plaid types as vector diagrams (left) and as the sum of sine-wave components (right). The aperture problem renders each component motion directionally ambiguous so that any vector terminating on the dashed "constraint" line would produce identical percepts; however, the intersection of constraint lines uniquely solves plaid direction. Adelson and Movshon (1982) proposed a neural implementation of this solution to account for perceived plaid direction. Alternatively, the "blob" features produced by grating luminance summation could determine perceived plaid motion, as their motion is identical to the IOC resultant. Neither proposal predicts the bias toward component directions in perceived direction of type II plaids. Another feature produced by summation of grating luminances, more evident in type II plaids, are second-order (contrast-defined) contours. These are the oblique gray contours seen in the figure which would move down and to the right. Plaid motion aftereffects (MAEs) reflect the adaptation of both component motions and both kinds of pattern feature. For both type I and type II plaids, the direction of the MAE is opposite to the perceived direction of motion.

type I and II plaids (Alais, Wenderoth, et al., 1994, 1997). Von Grünau and Dubé (1992) measured the strength of the MAE generated by plaids as a function of the difference in direction between the adapting plaid and the test stimulus (a counterphasing grating in which the relative contrast of the two opposing components could be adjusted to cancel the MAE). The tuning curves of MAEs obtained using plaid and grating adaptation stimuli were very similar. However, there was no evidence for stronger MAEs at test directions opposite to the plaid's component directions, further supporting the conclusion that MAEs (at least those measured with counterphasing test stimuli) might reflect adaptation at extrastriate locations.

5.4.3 *The Effect of Blobs*

When two drifting gratings with different orientations are superimposed to form a plaid, an array of bloblike features is formed by luminance summation at the intersections. These features always move in the true plaid direction because they are part of the plaid pattern itself, so detection of their motion would provide an alternative to the proposed IOC solution to plaid direction. Such a mechanism was originally rejected by Adelson and Movshon (1982) because 1-D dynamic noise disrupted coherence most when it masked the component directions rather than the plaid direction. Neither IOCs nor blobs can explain the component bias in the perceived direction of type II plaids (Ferrera and Wilson, 1990). Nevertheless, a good deal of evidence suggests that a mechanism sensitive to the motion of blob features does contribute to perceived plaid direction (Alais et al., 1997; Derrington and Badcock, 1992; Gorea and Lorenceau, 1991). Burke and Wenderoth (1993a), investigating earlier data from Wenderoth et al. (1988), compared plaid MAEs produced by grating components presented simultaneously with those produced by grating components presented in alternation (the latter condition precludes the formation of blobs). They suggested that the longer MAE duration obtained after simultaneous presentation could be due to the participation of a blob-tracking mechanism. In interocular transfer conditions (adapt using one eye and test using the other), the duration of the MAE following simultaneous adaptation fell to the same duration as the MAE following alternate adaptation (Burke and Wenderoth, 1993a, using type I plaids; Alais et al., 1994, using type II plaids). This result suggests that the blob-tracking mechanism is monocular. Importantly, MAE direction was more nearly opposite to the true direction of plaid motion in the simultaneous adaptation condition than in the alternating condition, as would be expected if the blob-tracking mechanism provided a veridical input to perceived plaid direction.

Alais et al. (1994) showed that the effectiveness of blobs depends on the spatiotemporal properties of the stimulus. Blobs formed by component gratings of high spatial frequency, low temporal frequency, and high contrast maximized the difference between alternately and simultaneously adapted MAEs, suggesting that these conditions optimally activate the mechanism responsive to the motion of the blobs. The inverse conditions minimized this difference. Using type II plaids, Alais et al. (1994) compared perceived direction for optimal and nonoptimal plaids, and found large differences in perceived direction: optimal blob plaids were perceived more veridically (by nearly 20 degrees) than nonoptimal blob plaids, largely overcoming the component bias known to occur with type II plaids (Ferrera and Wilson, 1990). Alais et al. (1997) further showed that simply increasing the spatial frequency of a type II plaid's components to about 6 cycles/degree progressively reduces the component bias. As would be expected if the blob-tracking mechanism were contributing substantially to perceived plaid direction, it slowly adapts during extended exposure to a drifting plaid, such that perceived direction is most veridical during the initial seconds of presentation (Wenderoth, Alais, et al., 1994).

This effect was examined more closely by Alais, Burke, et al. (1996a), who measured the perceived direction of type II plaids before and after monocular adaptation to the plaid's motion. When tested using the same eye, perceived direction differed significantly from perceived direction before adaptation, and was less veridical by an average of 7 degrees, showing more of the component bias. This difference was not found under interocular transfer conditions, suggesting the involvement of a low-level, monocular mechanism rather than a highly binocular structure such as area MT. The MAEs resulting from this period of adaptation appeared to move opposite to the postadaptation apparent direction, rather than the more veridical initial judgments of perceived plaid direction. The change in perceived plaid direction during a 30-second adaptation period reported by Alais et al. (1996) is in the opposite direction to that reported by Yo and Wilson (1992), who found that, over much briefer periods (on the order of 100 ms) after the onset of motion, plaids become *more* veridically perceived, moving from the component-biased (vector-sum) direction to the predicted IOC direction.

Overall, considerable evidence suggests that a monocular motion mechanism sensitive to the motion of a plaid's so-called blob features is involved in the perceived direction of plaids and their MAEs. In addition, adaptation studies have revealed that this same mechanism largely determines plaid coherence (Burke, Alais, et al., 1994; Alais, van der Smagt, et al., 1996b). Most of this evidence has emerged quite recently, and is interesting in light of the early history of plaid stimuli. It was generally acknowledged at the time that the ambiguous low-level input to the

motion system (the well-known aperture problem) posed a considerable challenge to models of motion perception, and seemed to require higher-level processes in order for object motion to be recovered. Plaid stimuli seemed to be the ideal tool for probing these higher-level motion processes since they are perceived to move rigidly (when coherence conditions are met) with an unambiguous direction despite being composed of directionally ambiguous components. However, it eventually became clear that the summation of component luminances produces second-order motion components and moving blob features which also make contributions to plaid motion perception. It is ironic that the mechanism detecting the motion of these features seems to be a low-level and monocular one, given that plaids were intended to probe higher-level motion processes. The importance of blobs is not surprising in view of their subjective salience when viewing a coherently drifting plaid. Component-related factors such as differences in component contrast (Adelson and Movshon, 1982), spatial frequency (A. T. Smith, 1992), and color (Krauskopf and Farell, 1990) also influence plaid coherence, but they may act as segmentation cues, with larger component differences along these dimensions increasing the likelihood that the components will be coded as independent.

5.4.4 Motion Aftereffects for Motion Components Separated in Depth

Verstraten et al. (1994d) demonstrated that under some circumstances, MAEs of bidirectional motion can indeed be bidirectional. Two simultaneous MAEs in different directions were sometimes observed with bivectorial adapting patterns, where the two random-dot patterns were located at crossed and uncrossed disparities, and with stationary test patterns also located at the same disparities. Oppositely moving adaptation stimuli showed significantly more instances of bivectorial MAE than orthogonally moving adaptation stimuli, where a unidirectional MAE was only seen in about 50 percent of the trials. These studies show that integration of motion directions prior to MAE generation is a very powerful mechanism which can persist even when adaptation and test stimuli are presented at different depth planes. On the other hand, the fact that bivectorial MAEs occur when tied to different disparities, but not when tied to color, for instance, speaks to the close interrelationship between channels responsible for direction analysis and for depth segregation (Kwas, von Grünau, et al., 1995, submitted; Trueswell and Hayhoe, 1993; Vallortigara and Bressan, 1991; see also chapter 4).

5.4.5 Attention and Bidirectional Motion Aftereffects

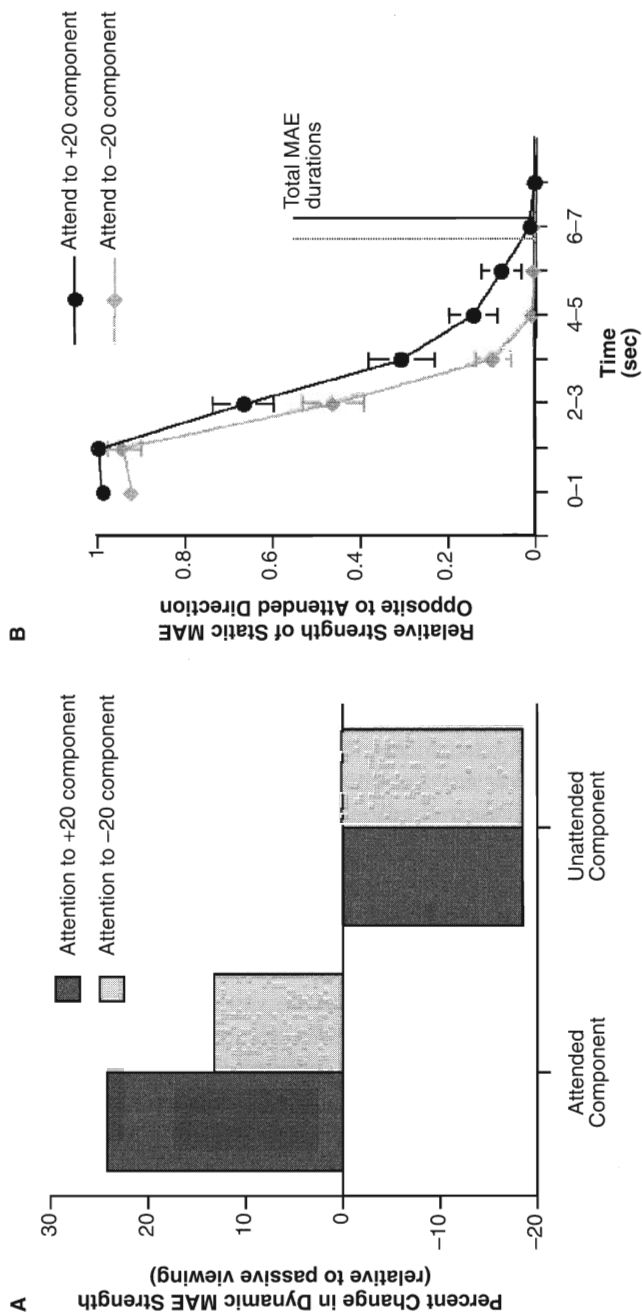
Given that attention to one component of a bidirectional random-dot pattern yields MAEs in the opposite direction (Lankheet and Verstraten,

1995), we may also wonder what would occur with attention to one of the two components forming a transparent plaid stimulus. Bertone, von Grünau, et al. (1997) had subjects adapt to a plaid, either while passively viewing it or while attending to one of the two components. MAEs were tested with counterphase flickering gratings as well as a stationary test plaid. When subjects attended to one of the transparent component gratings during adaptation, they perceived it as in front and dominant for 80 percent of the time. Relative to passive viewing, the strength of the flicker MAE increased when the test orientation matched the attended component and decreased when it matched the unattended component, as shown in figure 5.5A. The facilitatory and inhibitory effects were similar in size. With a stationary test plaid (figure 5.5B), illusory motion was usually perceived in the direction opposite to the adaptation plaid direction. When attention was directed toward one of the transparent component gratings during adaptation, the test plaid appeared to move opposite to the attended direction for a short time. As shown in figure 5.5B, this decayed rapidly and was replaced by motion opposite to the plaid direction. The overall MAE duration was not affected. Watanabe (1995) has also shown that when subjects attended to one edge of a series of translating pie-shaped wedges, tuning curves for static MAEs were shifted toward the orientation of the attended edge. Similar attentional influences have been observed with a bivectorial stimulus that could be seen as first-order motion in one direction or second-order motion in another (Iordanova, Riscaldino, et al., 1996). Again, flicker MAEs were stronger for the attended direction and weaker for the unattended direction, when compared to passive viewing. Taken together, the results of these studies indicate that attention both facilitates and inhibits motion mechanisms at or before the level of the MAE. The effects of attention on the MAEs of bivectorial motion suggest that component MAEs do exist and can be made "visible," but that without this attentional boost the bivectorial MAE is unidirectional.

In conclusion, MAEs to bivectorial motion do not appear bivectorial, even though the existence of the component MAEs can be demonstrated under conditions of selective attention or when direction is disparity-contingent. Most of the evidence suggests that the bivectorial MAE is produced only after the integration of the component motions. It is not clear, however, whether the same mechanisms contribute similarly to the formation of the integrated motion direction and the formation of the MAE direction.

5.5 Effects of Motion Adaptation on Global Motion Sensitivity

We now diverge from studies of the MAE to look at the effects of adaptation on global motion sensitivity, in order to compare the results

**Figure 5.5**

Attentional modulation of motion aftereffect (MAE) to a transparent plaid stimulus. A, Test stimuli were phase-reversing gratings with the same orientation as either the attended or unattended component of a moving plaid viewed during adaptation. For the attended orientation, MAE strength increased relative to a passive viewing baseline. For the unattended orientation, it decreased by a similar amount. B, At the beginning of a test period with a stationary plaid, the MAE was characterized by illusory motion opposite to the attended direction. The decay of this percept for successive 1-second intervals is indicated for one observer. Vertical bars represent ± 1 SE.

obtained using the two paradigms. Raymond and her colleagues conducted a series of experiments in which the adapting and test patterns were moving random-dot kinematograms (RDKs)—short-range apparent motion sequences consisting of successive presentations of random-dot displays in which a percentage of dots were displaced by a constant amount and in a single direction (signal dots) while remaining (noise) dots were displaced in random directions. In most of these experiments, the adaptation stimulus was 100 percent coherent and the percent coherence in the test was varied from trial to trial. Sensitivity to global motion in the test was assessed by determining the minimum percentage of signal dots needed for just accurate identification of signal direction. This measure is called the motion coherence threshold. We assume that the perception of global coherence is mediated by high-level motion-sensitive cells, probably in area MT, that integrate local motion signals over larger regions of space and are able to “smooth” across local differences in motion information to yield global signals of direction and velocity.

In the alert monkey, cells in area MT vary their response rates depending on the percentage of signal dots in RDK displays, and the motion coherence threshold of these monkeys can be manipulated by stimulation of these neurons (Salzman, Britten, et al., 1990b). Monkeys with lesions in area MT (Newsome and Paré, 1988) or humans with lesions in V5 (Zeki, 1993), an area homologous to MT (e.g., Baker, Hess, et al., 1991; Barton, Sharpe, et al., 1995), cannot perform motion coherence tasks normally. Raymond (1993b) asked observers to view a rightward moving 100 percent coherent RDK for 90 seconds and then alternated brief (196 ms) test stimuli with 5-second duration “top-up” intervals of 100 percent coherent motion to maintain adaptation. The task was to judge the global test direction from four possibilities: upward, downward, leftward, or rightward. A control condition in which subjects viewed a stationary dot display was used to determine the unadapted baseline threshold. Motion adaptation elevated motion coherence threshold when the test direction matched that of the adaptation stimulus. Group mean coherence threshold for five observers was elevated from a mean baseline 16 percent to 63 percent. This loss in global sensitivity was directionally selective because coherence thresholds for test stimuli with movement orthogonal or opposite to the adapting direction remained unchanged from baseline levels. Raymond (1993b) reported that the loss in sensitivity decreased as the difference in direction between adapt and test increased. The desensitization effect reached half its maximum when the directions differed by +35 degrees, thus describing the directional tuning of global analyzers.

This result is consistent with tuning functions found for MT cells in monkey (Dubner and Zeki, 1971; Felleman and Kaas, 1984) and is notably broader than the +20-degree bandwidth of tuning functions for V1

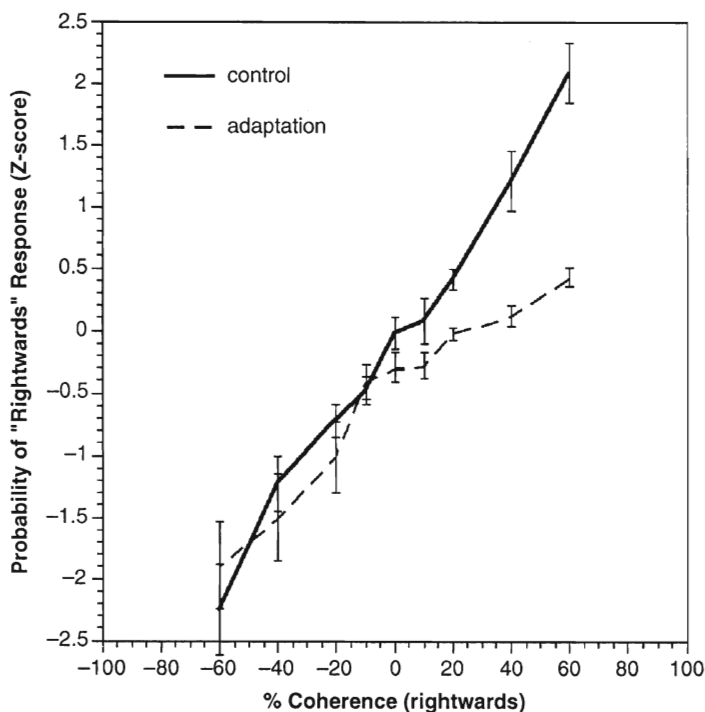


Figure 5.6

The group mean ($N = 5$) probability of a "rightward" response as a function of the percent coherence of rightward motion in the test stimulus after adaptation to rightward motion (dashed line) or to a static dot display (solid line). Negative percent coherence indicates leftward motion. Vertical bars represent ± 1 SE.

direction-selective cells (DeValois, Yund, et al., 1982; Schiller, Finlay, et al., 1976). The effect also showed broad speed selectivity (Raymond, 1994). Raymond and Braddick (1996) conducted a two-alternative forced choice ("rightward" vs. "leftward") experiment in which the psychometric function relating probability of a "rightward" response to percent coherence (varying coherence from 100 percent leftward through 0 to 100 percent rightward) was measured with and without prior adaptation to rightward motion. After adaptation, the slope of the function for rightward motion was flattened (figure 5.6) but the function for leftward motion remained unchanged, indicating that the mechanism mediating leftward detection was unaffected by rightward adaptation. These data suggest that adaptation measured in terms of changes in global motion thresholds affects high-level motion processing after a stage at which opponent interactions might occur (see chapter 7). Consistent with this interpretation was

the observation that if the adaptation stimulus consisted of 50 percent leftward and 50 percent rightward dots, that is, a bivectorial transparent display, coherence thresholds for test stimuli moving either leftward or rightward were elevated by an equal amount, while thresholds for the orthogonal directions were unaffected (Raymond, 1993b).

Many adaptation effects, including the MAE (e.g., Keck et al., 1976), increase in size as the duration or strength of the adapting stimulus is increased, and are not seen with brief or weak adapting stimuli. Raymond and Isaak (1998) tested whether the loss in the ability to see coherent motion in partially coherent displays reflected such a process. Long adaptation intervals and top-up intervals were replaced by discrete trials, each consisting of a successive adapt-test pair, with a long intertrial interval. The adaptation stimulus was either 100 percent coherent, rightward or leftward, or 0 percent coherent (i.e., motion noise). Test coherence thresholds were obtained for rightward and leftward motion. The remarkable result was that if the adapting stimulus was even as brief as 200 ms, test thresholds were significantly elevated when the adapt and test directions were the same. When the duration of the adaptation was varied from 200 to 2000 ms, the magnitude of threshold elevation increased modestly, reaching levels consistent with those found in the long adaptation experiment. Even with extremely brief adapting stimuli, motion coherence thresholds were elevated to nearly the same levels as found with much longer adaptation, suggesting that the neural processes leading to the adaptation effect occur very quickly and do not require prolonged exposure. In another experiment the coherence of the adapting stimulus was varied from 100 percent down to thresholds levels. As with stimulus duration, there was a modest increase in threshold elevation with increasing coherence, but even so, minimally suprathreshold coherence in the adapting stimulus raised thresholds significantly. These findings are inconsistent with an account of adaptation in terms of the gradual buildup of neural fatigue (theories of adaptation are discussed further in chapter 7).

Raymond, O'Donnell, et al. (1997) devised an experiment that was modeled after an attentional effect known as negative priming (e.g., Tipper and Cranston, 1985), but used the same stimuli as in the successive motion experiments described above. Results indicated clearly that selective attention modulates visual global motion perception. Attention to one direction caused a loss in sensitivity to that direction in a subsequent event, whereas actively ignoring a direction primed sensitivity for that direction. These effects are opposite in direction to more typical negative priming effects (where ignoring an item causes reduced subsequent processing). However, if we consider that changes in motion are more salient than continuation along a trajectory, mechanisms to reduce sensitivity to already coded information make sense. Another study, by Raymond and

O'Donnell (1996), further suggests that sensitivity reductions may only occur when motion stimuli are perceived as distinct objects. The authors first demonstrated threshold elevation effects when a brief (160 ms) passively viewed unidirectional RDK preceded a test RDK after an ISI (256 ms). In a second condition, the two motion stimuli were linked by a global object motion as the overall shape of the RDK changed between presentations ("morphing motion"; Tse and Cavanagh, in press). Rather than seeing the prime and test as discrete items, subjects perceived them as belonging to the same changing surface. Remarkably, threshold elevation effects disappeared. Given that the low-level stimulus properties were similar between the two conditions, these results suggest that motion adaptation effects depend on the perception of two different "objects."

In sum, the evidence reviewed here suggests that MAEs are but one instance of successive stimulus interactions which display a range of high-level properties. These experiments have indicated that much of the effect of motion adaptation on high-level motion perception does not depend on prolonged adaptation and is unlikely to be due solely to neural fatigue, as has been previously thought. By modifying classic psychophysical adaptation procedures to match widely used attentional negative-priming procedures, it has been demonstrated that successive stimulus interactions in motion perception are modulated by attention.

5.6 Conclusions

The past decade of research on the MAE has generated a new interest in a variety of higher-order effects that influence the illusion and has produced a number of surprising results. In particular, these new directions question the assumption that the MAE is a unitary preattentive effect generated by passive viewing of only unidirectional motion energy. Rather, it now appears that MAEs may be influenced by processing which spans a range of levels, both anatomically and functionally, and that MAE phenomena are the result of complex interactions between stimuli and tasks.

These complex interactions within MAE phenomena provide a rich domain for exploring motion processing and its interactions with other *perceptual and cognitive processes*. As we have observed, complex stimuli and tasks can be used to probe the MAE. For example, interocular transfer studies and higher-order motion stimuli have suggested that static MAEs may occur at earlier stages than dynamic MAEs. However, not only can complex factors be used to understand the MAE but the MAE can also be used to examine higher-order aspects of stimulus coding and the attentional state. Examples include the use of MAEs to probe the stages of motion integration (for both plaid components and random dot

stimuli) and the interactions between attention, priming, and motion processing.

From this recent literature, three primary themes and future research directions emerge. The first concerns the issue of whether static and dynamic test patterns reflect common underlying mechanisms or tap separate levels of motion processing. The list of differences between static and dynamic stimuli is certainly large and growing (see table 5.1), and is difficult to reconcile with a single mechanism. However, a number of studies suggest similarities in the properties of the two test types (e.g., Nishida et al., 1997a) or provide models which can account for some of the data without the need for distinct mechanisms (Wainwright and Cavanagh, 1997). Nevertheless, other properties, such as the differences in interocular transfer, pose difficulties for a single-mechanism account. One possible resolution is that there is a continuum of properties which are differentially tapped by the temporal frequency ranges typically used for static and flicker tests. A second possibility is that multiple factors may account for the properties of flicker MAE tests—a low-level motion adaptation factor which is not fully binocular but is dominant in peripheral vision, and a second high-level factor which is closely linked with attention, contributes to apparent motion correspondence matches, and is more dominant in central vision.

Second, the MAE depends on more than just the first-order motion components making up a stimulus. Second-order motion displays and attentional shifts can also generate MAEs, at least for dynamic tests. Further research will be useful in clarifying interactions between first-order, second-order, and “third-order” (i.e., attention-based) motion systems. Furthermore, even within first-order motion, MAEs are specific to the integrated direction of multiple components rather than the directions of the components themselves (with one intriguing exception in the case of components separated in depth). This research has provided evidence both for low-level monocular mechanisms involved in “blob-tracking” and for higher-level, possibly extrastriate, integration of components before the site where the MAE is generated. Again, these results raise interesting issues concerning the relationships between these mechanisms, and their physiologic substrates.

Third, the MAE may be much more susceptible to cognitive factors, particularly attention, than previously believed. Many recent studies show that responses to motion stimuli and their components can be enhanced or suppressed depending on whether they are attended and selected or not. Furthermore, attention may act in multiple ways, both modulating the strength of motion signals and possibly mediating correspondences in a motion system capable of producing distinct aftereffects. Finally, motion sensitivity measures suggest that only brief, suprathreshold motion

instances may be necessary to produce "adaptation" effects and that such effects may be related to other cognitive interactions between successive stimuli such as priming.

In sum, the main theme that emerges from this chapter is that there may not be a single unitary MAE measured by all studies which address it. Rather, the properties of the MAE may depend substantially on the type of adaptation stimulus employed (first- or second-order, components or integrated patterns), the type of test stimulus used (static or dynamic), and the subject's state during adaptation (including attentional and response selection factors). Many of these data are preliminary, the theories are sometimes speculative, and, as might be obvious from some of the earlier chapters in this book, a number of different interpretations are available. Nonetheless, these different theories and speculations provide an excellent impetus for further research into higher-order factors in the MAE.