

Integration of Speed Signals in the Direction of Motion

Dawn Vreven and Preeti Verghese

Smith-Kettlewell Eye Research Institute

2318 Fillmore Street

San Francisco, CA 94115

Telephone: (415)345-2097

E-mail: dawn@ski.org

ABSTRACT

Speed discrimination tasks were used to examine the spatial and temporal characteristics of the integration mechanism involved when signals are extended in the direction of motion. We varied the aspect ratio of a signal patch whose speed differed from the background while holding the area of the signal patch constant, such that the signal patch could be either extended in the direction of motion or extended orthogonal to the direction of motion. Speed discrimination thresholds decreased dramatically as the signal patch was extended in the direction of motion.

The spatial and temporal integration regions were larger than would be expected if the integration mechanism were a low-level motion detector. The mechanism was tuned for direction of motion. The data are discussed with reference to two alternative integration mechanisms: a low-level detector that is elongated in the direction of motion, and a higher-level integration mechanism characterized by cooperative or facilitory interactions between low-level detectors tuned to the same direction of motion. Our data are consistent with a second-level, direction-specific process that integrates the responses of low-level of motion detectors.

INTRODUCTION

It is well accepted that motion is initially encoded in the brain by local motion detectors. Energy models of this initial stage of motion detection assume that an area of visual space is tiled with directionally-specific motion detectors at multiple scales (e.g., Watson & Ahumada, 1985). Psychophysical evidence has suggested that these low-level motion detectors have circular receptive fields (Anderson & Burr, 1991; Watson & Turano, 1995; Georgeson & Scott-Samuel, 2000) and integration times of between 80-120 ms (Lappin & Bell, 1976; Watson, 1979).

This description of motion processing does not take into account the fact that an extended motion sequence in a consistent direction is a special stimulus for the visual system. There is a great deal of evidence that human observers can detect signals extended in the direction of motion more accurately than predicted by independent low-level motion detectors (Anstis & Ramachandran, 1987; Snowden & Braddick, 1989; Verghese, Watamaniuk, McKee & Grzywacz, 1999; Verghese, McKee & Grzywacz, 2000). Watamaniuk, McKee & Grzywacz (1995), for example, found that one signal dot moving in a consistent direction amid hundreds of random-direction noise dots was easily detectable, even though on a frame-by-frame basis the motion of the signal dot was indistinguishable from a noise dot. Studies of motion coherence also reveal an advantage for signals extended in the direction of motion (van Doorn & Koenderink, 1984; Fredericksen, Verstraten, & van de Grind, 1994). This work measured the luminance signal-to-noise ratio (SNR) required to detect the direction of horizontally translating coherent dots in a display of incoherent noise. van Doorn & Koenderink (1984) manipulated the region in which coherent dots appeared. In one condition, the

coherent dots appeared in a narrow vertical strip; van Doorn and Koenderink increased the horizontal extent in which coherent motion could appear, while keeping the vertical extent and the coherent dot density constant. Luminance SNRs decreased as the coherent dot region was extended in the direction of motion.

The goal of work in this area is to determine how such performance could arise given the properties of low-level motion detectors. There is general agreement that combining motion signals over space and time could explain the data, but there are several ways that such spatiotemporal integration could occur and discerning among the alternatives is difficult. Two of the most straightforward explanations are noise-limited probability summation and multiple scale recruitment. van Doorn & Koenderink (1984) proposed both kinds of motion integration to explain their results. First, they postulated that as the area of the coherent region grew, a greater number of detectors received a signal. This is the notion of probability summation, in which increasing the number of activated detectors yields increased performance. Secondly, they postulated that detectors at larger spatial scales became available as the signal was elongated in the direction of motion. A thin strip of coherent motion, for example, optimally activates small detectors. Larger detectors may receive some coherence signal, but the signal is not optimal for the size of the detector. As the coherent region becomes extended in the direction of motion, however, both the smaller detectors and larger detectors receive a strong coherence signal, and so both detector sizes can contribute to increased performance.

One reason that discerning among integration schemes is so difficult is because it is difficult to isolate only one integration mechanism. Fredericksen et al. (1994) attempted to simplify the problem by ruling out detectors at multiple scales. They used

a paradigm similar to van Doorn & Koenderink's, but used dots with 1 frame lifetimes. This strategy isolates a population of motion detectors sensitive to only a single spatial displacement, and thus reduces the possibility of recruiting detectors at other scales. Fredericksen et al. found, as before, that luminance SNRs improved as the region of coherent dots was extended in the direction of motion. The improvement they found, however, was greater than could be predicted based on probability summation alone. They argued that probability summation among detectors spatially *elongated* in the direction of motion could explain the data.

An alternate approach to understanding motion integration involves holding the area of the motion signal constant while changing its aspect ratio (Todd & Norman, 1995). This simple manipulation, depicted in Figure 1, changes the orientation of a rectangular signal patch relative to the direction of motion. There are a number of advantages to using this stimulus. One advantage is that if motion detectors are circular, then a similar number of detectors will be activated for a signal patch of a given area, regardless of the shape of the patch. Holding the area of the signal patch constant therefore makes the probability summation explanation less likely. A second advantage of holding the signal area constant while changing the aspect ratio of the patch is that the use of multiple scales of detectors becomes less likely. This is shown schematically in Figure 1A. Assuming that the detectors are circular and are smaller than the stimulus envelope, there is no advantage for a stimulus elongated in the direction of

Insert Figure 1 about here

motion. Thus, the aspect ratio manipulation eliminates simple noise-limited probability summation as an integration mechanism and is designed to activate motion detectors of a single scale.

Two remaining integration mechanisms that could be involved in our task are shown schematically in Figure 1. The first mechanism is a detector elongated in the direction of motion, represented by an ellipse in Figure 1B. The ellipse represents the spatial and temporal neighborhood where signals are pooled. Improvement in performance is due to increasing amounts of signal falling within the detector's receptive field in a process known as linear or physiological summation (as opposed to noise-limited probability summation). If the signal patch is oriented *perpendicular* to the direction of motion (left), for example, the elongated detector will combine motion information in the signal patch with background noise. The signal-to-noise ratio would be low and as a result, motion detection or discrimination should be difficult. Now consider how performance will change as the shape of the signal patch changes from vertical (perpendicular to the direction of motion) to square. As the horizontal dimension of the signal patch grows parallel to the long extent of the elongated detector, it receives an increasingly better signal. The signal-to-noise ratio should increase as the stimulus is extended in the direction of motion, leading to improved detection and discrimination performance. In this case, thresholds plotted as a function of aspect ratio will decrease steadily. When the signal patch is further extended so that it matches the length of the detector, the detector integrates motion information from the signal patch alone and performance improves further.

A number of investigators have proposed a second kind of integration mechanism that relies on cooperativity among motion detectors (Chang & Julesz, 1984; Williams,

Philips & Sekuler, 1986; Yuille & Grzywacz, 1988; Snowden & Braddick, 1989; Nawrot & Sekuler, 1990; Gurney & Wright, 1996; Grzywacz, Watamaniuk & McKee, 1995). This mechanism is schematized in Figure 1C. The idea is that low-level motion detectors have connections with other detectors in a spatial and temporal “neighborhood”. Only detectors that are tuned to similar stimulus properties, such as preferred direction, interact with one another. When a single detector is activated by its preferred stimulus, it in turn activates those detectors to which it is connected. Activation from neighboring detectors with similar tuning, combined with activation from the stimulus, accounts for the enhanced performance for signals with consistent direction. This mechanism is “second stage” in the sense that it refers to interactions between the responses of low-level motion detectors.

We explored motion integration by looking for the presence or absence of stimulus configuration effects. If changing the aspect ratio of the signal patch improves thresholds, then integration in the direction of motion has occurred. Changes in the rate of threshold improvement can also yield information about the extent of motion integration. As the stimulus is extended beyond the boundaries of the integration mechanism, for example, threshold improvement should continue at a slower rate or perhaps become asymptotic. Our purpose was to 1) confirm that the aspect ratio manipulation yields the same kind of results that have previously been found and 2) determine which integration mechanism (or combinations of mechanisms) best account for the data.

EXPERIMENT 1: COHERENCE DETECTION

Experiment 1 is a bridge between the work of van Doorn & Koenderink (1984) and Fredricksen et al. (1994), which measured luminance signal-to-noise ratios to detect the direction of motion of a coherent patch, and our paradigm, which keeps stimulus area constant while changing aspect ratio. Does sensitivity to the coherent patch increase as it is extended in the direction of motion?

METHOD

Observers

A total of six observers participated in the experiments reported in this paper. Not all observers were available to serve in all experiments. Each observer had normal or corrected to normal vision. Two observers were the authors, and the remaining four were naïve about the purpose of the experiments. Three observers participated in Experiment 1, one author and two naïve observers.

Stimuli and Experimental Procedure

Stimuli were generated using custom software and displayed on a 38.1 cm CRT monitor with a P4 phosphor and a refresh rate of 71 Hz. Stimulus displays were dynamic random dot cinematograms, viewed from 100 cm. Individual bright dots, subtending 3 min^2 at the 1 m viewing distance, could appear anywhere within a circular display aperture whose diameter was 12.6 deg arc (area = 124.7 deg^2). Typically 512 dots appeared in the aperture, yielding a dot density of 4.1 dots/deg^2 . The luminance of the dots, measured by a Pritchard photometer on a 6 min by 6 min test patch, was 130

candelas/m² at the 1 m viewing distance; the background luminance was 45 candelas/m².

A "windowing" technique was used to generate the cinematograms, as illustrated in Figure 2. Two windows were presented simultaneously for each stimulus display --

Insert Figure 2 about here

a rectangular "signal" window surrounded by a circular "background" window. The density of dots in each window was equivalent. The background window was the size of the display aperture (minus the size of the signal window) and contained dots that were randomly re-plotted on every frame. The signal window was centered in the display aperture with an area of 9 deg². The shape of the signal window varied depending on the signal aspect ratio being tested (see below). Displays were presented in a two-interval forced choice paradigm. On each trial, a fixation point appeared in the center of the screen. After a brief delay (approximately 300 ms) the first of two stimulus intervals was presented. Each interval lasted for 200 ms (14 frames), with a 350 ms blank screen presented between each interval. Two types of motion were generated for dots in the signal window, corresponding to the signal and noise intervals. For the noise interval, all dots within the signal window were randomly-replotted as in the background window. For the signal interval, a proportion of randomly-chosen dots translated coherently (all leftward or all rightward) at 10 deg/sec, while the remaining dots were randomly re-plotted. The observers' task was to indicate which interval contained coherently moving dots. The lifetime of coherent dots was limited to two frames to prevent observers from tracking any particular dot. After two frames of motion, a

coherent dot was replotted at a random location somewhere within the circular aperture. The proportion of coherent dots typically varied from around .10 to .30 in five equal steps; this range was changed for individual observers if more or less coherence was required to determine the coherence threshold. Finally, note that because the signal window had an area of 9 deg^2 and was presented at a density of 4 dots/deg^2 , it contained 37 dots per frame on average. Therefore, a coherence proportion of .25 corresponds to 9 coherent dots per frame appearing within the signal window.

The shape of the signal patch in which the coherently-moving dots could appear was the manipulation of interest. The area was held constant at 9 deg^2 , while the height and width of the patch were manipulated. The center of the signal patch was always coincident with the center of the screen. The shape of the signal patch is expressed in terms of horizontal aspect ratio (AR), which equals width/height. At one extreme, the width of the patch was set to 1 deg and its height was set at 9 deg; this vertically-oriented patch had an aspect ratio of .11. In other words, a dot translating at 10 deg/sec covered 1 deg in 100 ms, which was designed to correspond to the integration time of a local motion detector. At the other extreme, the signal patch was given a width of 9 degrees and a height of 1 deg, yielding an aspect ratio of 9. Seven aspect ratios were used: .11 (a vertical rectangle 1 deg wide), .165, .33, 1 (a 3 deg-by-3 deg square), 3, 6 and 9 (a horizontal rectangle 9 deg wide).

One session consisted of 100 trials of a single randomly-selected aspect ratio (20 trials x 5 coherence levels). A coherence threshold was then determined for a particular aspect ratio by measuring the proportion correct at each level of coherence. This data was fit with a Weibull function of the form

$$p = 1 - 0.5 \exp [(-x/\alpha)^\beta]$$

where p is the probability of a correct response, x is the proportion coherence, α is the threshold proportion coherence associated with 82% correct performance, and β is the slope parameter. Each observer repeated this procedure a total of three times for each of the 7 aspect ratios. Thus, each threshold in the figures below is based on the average of 3 separate threshold estimates, or a total of 300 observations.

Results

 Insert Figure 3 about here

Figure 3 plots the coherence thresholds and standard errors as a function of aspect ratio for three observers. In general, thresholds decrease by 50% as the signal patch becomes extended in the direction of motion. To test whether thresholds significantly improve as the aspect ratio changes, a two-factor (AR[7] x Session[3]) repeated-measures ANOVA was performed on the coherence threshold data. The effect of aspect ratio was significant, $F(6,12) = 7.11, p = .002$. At the largest aspect ratios, coherence thresholds are within the range of previous reports for both human and primate observers (Newsome & Pare, 1988; Britten, Shadlen, Newsome, & Movshon, 1992). This is true despite the fact that dot density in the current experiment was much lower than that used in previous studies (4 dots/deg² in the current experiment vs. 16.7 dots/deg² in Britten et al. [1992], for example). Further, the rate of improvement for all three observers decelerates as the aspect ratio increases.

Discussion

The observed improvement in coherence thresholds as the signal becomes extended in the direction of motion is consistent with previous studies. The improvement cannot easily be accounted for by probability summation or multiple detector scales. One additional observation suggests that threshold improvement is not due to physiological summation within a single elongated detector. The spatial extent of integration was larger than 3 deg for all observers, due to the fact that thresholds continue to improve beyond an aspect ratio of 1 (Figure 2). Assuming that the spatial and temporal extents of the integration mechanism are non-separable (i.e., speed tuned), a large spatial extent implies a long temporal extent. Thus, the integration mechanism has an integration time longer than the 100ms characteristic of low-level motion detectors (Watson, 1979; Fredericksen et al., 1994). Recall that coherent motion speed was 10 deg/sec; at this speed, 1 deg is traversed in 100 ms. The time constant of a mechanism tuned to this speed that spatially integrates for *over* a deg, therefore, must be greater than 100 ms. Thus, large spatial integration regions imply integration times longer than those typically associated with low-level motion detectors. In order to account for the data, the elongated detector would be required to have a large time constant.

EXPERIMENT 2: SPEED DISCRIMINATION

Because noise dots are randomly-replotted in the coherence paradigm, they contain wide speed and direction distributions. The signal dots, however, move in one

direction at a single speed. Thus, the coherence paradigm uses both dot direction and speed to segregate the signal from noise. We wondered about the relative importance of these two properties for motion integration. To explore this idea, we used background dots translating in a uniform direction at a speed which was distinct from the speed of the target dots. We manipulated the speed of the signal dots relative to the background dots in a discrimination paradigm. There is evidence that relatively imprecise local speed signals are combined over time (Watamaniuk & Duchon, 1992; Bravo & Watamaniuk, 1995; Snowden & Braddick, 1989, 1991), and we used this approach to further characterize the integration mechanism.

METHOD

Procedure

The two authors and three naïve observers participated in this experiment. The same equipment and windowing strategy involving background dots and signal patch dots was used to generate the stimuli. The aspect ratio of the signal patch was manipulated as before. Dot density for both windows was maintained at an average of 4.1 dots/deg². There was no constraint on dot lifetime.

Background dots always translated at 10 deg/sec. A background dot moving out of the viewable area of the circular aperture was re-plotted at a random location within the background window. A background dot moving into the signal patch window disappeared at the border between the two windows and was similarly re-plotted at a random location within the background window. Signal patch dots translating out of the signal window also disappeared as they approached the border of the window, and

were replotted at a random location within the signal window. Signal dots translated at one of five speeds that typically varied from 10 deg/sec to 12 deg/sec in five equal steps. The range of increment dot speed in the signal patch could be altered for individual observers if required to determine the speed discrimination threshold. Background and signal patch dots always translated in the same direction (randomly either rightward or leftward). Because extended lifetime dots were used, an implicit border was present between the signal patch and background windows within the stimulus display. This border was evident at the vertical contours of the signal patch and was created by the appearance and disappearance of both signal patch dots and background dots at the border between windows. The border appeared on both intervals of each trial and therefore could not be used to distinguish which interval contained the signal.

On each trial, the observer was presented with two, 200-ms intervals using the method of constant stimuli. In one interval, both background and foreground dots translated at 10 deg/sec; in the other interval, background dots continued to translate at 10 deg/sec while signal patch dots translated at a randomly-selected speed from the pool of five speeds. Each speed appeared equally often. The observers' task was to indicate which interval contained the faster target patch speed.

Results and Discussion

Insert Figure 4 about here

Speed discrimination thresholds appear in Figure 4. All observers show a decrease in speed threshold (increase in speed sensitivity) as the signal patch becomes

extended in the direction of motion. For most observers, thresholds for the largest aspect ratios were lower by a factor of two or more compared to thresholds for the smallest aspect ratios. As before, a two-factor (AR[7] x Session[3]) within-subjects ANOVA on the speed discrimination thresholds indicated that the effect of aspect ratio was significant ($F[6,24] = 2.70, p = .038$). Additionally, the pattern of threshold improvement is similar to that found in Experiment 1. Four of five observers have thresholds for the largest aspect ratios that are within the error of one another, indicating that performance has ceased to improve as the signal becomes extended in the direction of motion. This decelerating rate of improvement generally occurs at aspect ratios greater than 1, indicating that spatial integration in the direction of motion is greater than 3 degrees.

In sum, the speed discrimination experiment yielded results that were qualitatively similar to the coherence paradigm of Experiment 1. For small aspect ratios, the Weber fraction for speed (the smallest difference in speed that can be reliably judged; $\Delta\text{speed}/\text{speed}$) ranged from .10 to .20. As the aspect ratio increased and the signal became elongated in the direction of motion, the Weber fraction dropped to around .05. This value is within the .04 to .08 found as the smallest difference in speed that can be judged reliably by human observers (McKee, 1981; Watamaniuk & Duchon, 1992). Experiment 2 demonstrates that speed signals are integrated in the direction of motion much as coherence signals are. High speed discrimination thresholds result when the signal is restricted to small extents in the direction of motion. This implies that speed information from the signal and background is being combined. This could result from either an elongated low-level detector or from a second-stage integrator that combines the outputs from circular low-level detectors.

EXPERIMENT 3: DURATION

Experiments 1 and 2 were concerned with the spatial extent of the motion integration mechanism. Our purpose here is to determine the temporal extent of the mechanism. The strategy is to look for the characteristic improvement in performance at various stimulus durations. Consider, for example, the output of a motion mechanism to displays of 25 ms. Such a short duration does not maximally stimulate low-level motion detectors, whose time constant is around 100 ms (Watson, 1979; Fredericksen et al., 1994). In fact, for circular detectors, we would expect little change in performance with increasing aspect ratio. For elongated detectors, however, performance should improve as the signal more closely matches the shape of the detector. Now consider longer durations, which exceed the integration time of the low-level motion detectors. If speed discrimination is mediated by elongated low-level detectors, there should be no change in the pattern of thresholds as a function of aspect ratio once the integration time of these units is exceeded. Absolute thresholds might decrease with increasing duration, but the pattern of threshold improvement should be similar for all durations. If, however, performance depends upon interactions between these low-level motion detectors, thresholds should decrease as a function of aspect ratio only at longer durations. We are therefore interested in the duration at which speed discrimination thresholds are affected by stimulus configuration. The logic of the experiment is to determine the presence or absence of stimulus configuration effects at various durations, allowing an estimate of the extent of temporal integration.

Procedure

The two authors and two naïve observers participated in this experiment. The procedure and stimulus displays for this experiment were identical to those of Experiment 2. Stimulus durations of 50 ms, 100 ms and 400 ms were tested. Not all observers participated in all duration conditions. For comparison, the 200 ms duration data from Experiment 2 is also included.

Results and Discussion

Insert Figure 5 about here

Figure 5 plots speed discrimination thresholds as a function of aspect ratio. Each panel shows the data for one observer. Durations of 100 ms are shown with triangular symbols. For comparison, 200 ms duration data are plotted here with circular symbols. Square symbols show the data for 400 ms durations. Because we are interested in the presence or absence of configuration effects, each observer's data has been fit with a line. A line with zero slope indicates no stimulus configuration effect.

Consider first the data from observers DV, FP and PV. At 100 ms duration (triangular symbols) , these observers are unaffected by the aspect ratio manipulation. Each of these three observers' data is best fit with a line of near zero slope. The flat data functions for these observers suggest that the integration mechanism was not engaged at this brief duration. Next consider the 200 ms data (circular symbols) for the same three observers. In each case, a stimulus configuration effect is present in the form of decreasing thresholds as aspect ratio increases, with a decelerating rate of improvement

at large aspect ratios. Increasing the duration to 400 ms (square symbols) further improves performance, particularly at large aspect ratios. Speed discrimination thresholds from these three observers were submitted to a three-factor (Duration [3] x AR[7] x Session[3]) repeated-measures ANOVA. Significant main effects of duration and aspect ratio were found, $F(2,4) = 19.64, p = .009$ and $F(6,12) = 10.13, p = .0001$, respectively. More importantly, however, the interaction of duration and aspect ratio was significant, $F(12,24) = 3.72, p = .003$. Three post-hoc comparisons revealed that the threshold at 100 ms did not significantly improve from the smallest aspect ratio to the largest aspect ratio (mean thresholds of .167 and .147 for aspect ratios .11 and 9, respectively; $F[1,2] = 1.16$). Thresholds did significantly improve for the 200 ms duration (mean thresholds of .139 and .056 for aspect ratios .11 and 9, respectively; $F[1,2] = 46.53, p = .02$), and narrowly failed to achieve significance for the 400 ms duration (mean thresholds of .109 and .024 for aspect ratios .11 and 9, respectively; $F[1,2] = 12.32, p = .07$). Because stimulus configuration effects were present at 200 ms but not at 100 ms, we estimate the extent of *temporal* integration for these three observers to be between 100 ms and 200 ms.

One observer (NK) shows stimulus configuration effects at 100 ms duration. We wondered whether this observer would show a similar effect at a shorter duration. If configuration effects are due to an integration mechanism that operates on the responses of low-level motion detectors, there should be some duration at which low-level motion detectors are still engaged, but no configuration effect is present. To check this possibility, we tested this observer at 50 ms duration (diamond symbols). Thresholds in this case are virtually unaffected by aspect ratio. This outcome suggests that the observer has a slightly shorter low-level time constant than the others.

What do these results tell us about the integration mechanism? Three of four observers were unaffected by stimulus configuration at 100 ms, which is a typical time constant for a first-stage motion detector. This outcome suggests that the stimulus configuration effects we find are not mediated by first-stage detectors. Rather, we contend that it takes longer for the second-stage mechanism to integrate the outputs of low-level detectors. In other words, we propose that the integration mechanism responsible for stimulus configuration effects operates on the responses of first-stage motion detectors, and therefore requires longer than 100 ms to develop. What about the observer (NK) who shows some evidence of motion integration at the brief 100 ms duration? We attribute this to variation in the integration time of first-stage motion detectors. Integration times as low as 80 ms have been reported for first-stage motion detectors (Lappin & Bell, 1976). It is possible that this observer has a shorter low-level time constant and, therefore, a shorter extent of temporal integration.

An elongated motion detector mechanism cannot easily account for the data. For a given stimulus duration, an elongated detector would be better stimulated by a signal extended in the direction of motion due to physiological summation within the detector. Thus, an elongated detector should show threshold improvement even at 100 ms. The lack of stimulus configuration effect at 100 ms is inconsistent with this prediction.

EXPERIMENT 4: SPEED DISCRIMINATION WITH ORTHOGONAL DIRECTIONS OF MOTION

As noted above, the signal dots in the coherence experiment (Experiment 1) differed from the noise dots in both speed and direction. The high thresholds for small aspect ratios in Experiment 2 suggest that the integration mechanism combines speeds from the signal patch and the background. Our aim in this experiment was to test whether the mechanism also integrates over different directions of motion. We used a speed discrimination task identical to that of Experiment 2, with one difference: background dots moved orthogonally to dots in the signal patch. If the mechanism integrates over both speed and orthogonal directions of motion, then we expect the familiar outcome of initially decreasing thresholds followed by asymptotic performance. If, however, the mechanism does not integrate over speed and direction, thresholds should be relatively unaffected by the aspect ratio manipulation.

There is an additional motivation for this experiment. It could be argued that the improvement in speed thresholds as the signal becomes extended in the direction of motion is a result of poor speed estimates for the vertically-elongated ($AR = .11$) signal patch. Recall that the width of this signal patch was 1 degree, and the base speed was 10 degrees/second. Thus a signal dot would take 100 ms to traverse this path. Because target dots start at a random position within the signal window, each target dot traverses (on average) .5 degrees and is visible for only 50 ms of the 200 ms interval. This relatively brief presentation could negatively affect local speed estimates. Note that changing the direction of the background dots leaves all speed information in the signal patch intact, including the average 50 ms traversal time for the smallest aspect ratio. If the stimulus configuration effects we find are due to poor speed estimates for small aspect ratios, then we should expect the same pattern of threshold improvement

in the current experiment because the motion in the signal patch is identical in both cases.

Procedure

Experiment 4 was identical in all respects to Experiment 2, with the following exception: the background dots translated perpendicular (i.e., upward or downward) to the foreground dots, which translated randomly leftward or rightward. In this way, both the direction and the speed of the background dots differed from the signal patch dots. All other aspects of the experiment were identical to Experiment 2. Stimulus duration was 200 ms.

Results and Discussion

Insert Figure 6 about here

Thresholds for speed discrimination with orthogonal directions of motion are shown in Figure 6. These thresholds should be compared with those in Figure 4, where the same direction of motion appeared in the signal patch and the background.

Thresholds are relatively unaffected by stimulus configuration when signal patch and background dots move in orthogonal directions. A two-factor (AR[7] x Session[3]) ANOVA performed on the data indicated that the overall effect of aspect ratio was not significant ($F[6,24] = 1.35, p = .27$). Despite this fact, it appears that some observers (DV and NK) are less affected by orthogonal directions of motion than others. The size of the stimulus configuration effect can be calculated by forming a ratio of thresholds

from aspect ratios .11 and 9. These ratios appear in Table 1, along with the ratios for the same 5 observers in Experiment 2 (same direction of motion in target and background). A threshold ratio of 1 indicates no stimulus configuration effect; ratios greater than 1 indicate that thresholds decrease as aspect ratio increases. The difference between the size of the stimulus configuration effect for same vs. orthogonal directions of motion is small for observer DV (2.0 and 1.7, respectively), and observer NK shows no difference (1.8). Consider, however, that observer RC has a threshold ratio of 0.8 for orthogonal directions of motion, indicating that thresholds increase with increasing aspect ratio. This variability between observers probably reflects a limit in the resolution of our paradigm; if, for example, we had tested aspect ratios smaller than .11 or greater than 9, it is likely we would have found larger and/or more reliable differences between the outcomes of Experiments 2 and 4.

Insert Table 1 about here

In general, directional consistency (at least within the 90 degree direction change used here) is required to observe large stimulus configuration effects. When signal and background dots move in different directions, motion integration does not occur at the same rate. Rather, the data are consistent with an integration process that occurs between background and signal dots moving in the same direction.

The relative lack of threshold improvement as aspect ratio increased with orthogonal background/signal dot directions argues against the notion that signal patches with small aspect ratios are subject to poor speed estimates relative to signal patches with large aspect ratios. If this were the case, we would expect the same pattern

of threshold improvement in the current experiment as in Experiment 2, because the signal patch was identical. The stimulus configuration effects found in Experiment 2, therefore, can not be attributed to faulty local speed estimates independent of motion integration.

As before, it is unclear how an elongated motion detector could account for this outcome because such a detector would be better stimulated by a signal extended in the direction of motion. Interaction between local detectors, however, could account for the data quite readily because the facilitation only occurs among low-level detectors with similar speed and direction preferences.

GENERAL DISCUSSION

Our purpose was to explore the mechanism responsible for the integration of signals extended in the direction of motion. To this end, we used speed discrimination tasks to characterize the spatial and temporal extent of motion integration. We kept the area of a signal patch constant and found that thresholds improved as we increased the extent of the signal in the direction of motion. In general, thresholds continue to improve beyond an aspect ratio of 1, indicating spatial integration of at least 3 degrees. The idea of a large area of spatial integration is consistent with the work of Todd & Norman (1995), who found spatial integration extents of 10 to 20 degrees. Our data also indicate a lower bound for the time required for threshold improvement in the direction of motion. Most observers showed the characteristic pattern of threshold improvement with stimulus durations of 200 ms; one observer showed the effect at a duration of 100 ms, suggesting a very small temporal integration duration. The variability of our data

may not be surprising: Lappin & Bell (1976) show that the integration times of low level motion detectors can be as brief as 80 ms.

Our results, like those of van Doorn & Koenderink (1984), Fredericksen et al. (1994), and Todd & Norman (1995), show an increased sensitivity to signals extended in the direction of motion. Unlike some of this work, however, the aspect ratio manipulation maintained a constant signal area, and therefore a constant number of first-stage motion detectors were exposed to the stimulus. This manipulation alone rules out simple probability summation as an explanation for the motion integration we have observed, since such an explanation requires the stimulus to be presented to increasing numbers of detectors.

One important issue, pointed out by Todd & Norman (1995) and Fredericksen, Verstraten, & van de Grind (1997), is that of low spatiotemporal correlation for signal patches oriented perpendicular to the direction of motion (i.e., signal patches with small aspect ratios). Signal strength decreases as the frame-to-frame spatial displacement (i.e., "hop size") of dots increases relative to the size of the stimulus window. This is an inherent property of the stimulus and is a concern when the step size is large relative to the width of the patch in the direction of motion. The hop size for our base speed of 10 deg/sec was .14 deg per frame, which is small relative to the 1 deg wide stimulus window for aspect ratio = .11. Our hop size, therefore, is one-seventh the size of the stimulus window. In fact, the data of Fredericksen et al. show that performance is unaffected by low spatiotemporal correlation when the hop size is as large as one-fourth the size of the stimulus window. We also addressed concerns about differences in stimulus energy as a function of aspect ratio by conducting a simulation. We computed the output of a local motion energy detector to a target patch containing dots moving in the detector's

preferred direction (rightward). We matched the dot speed and dot density to those of our experimental stimuli, and we re-plotted dots that moved out of the target patch in a similar manner (see Methods section, Exp.2). The circularly-symmetric detector had a spatial extent of 1 deg, an integration time of 100 ms, and was tuned to rightward motion at a speed of 10 deg/sec. We simulated two conditions, corresponding to aspect ratios of .11 and 9. The spatial extent of 1 deg matched the width and height of the signal windows for aspect ratios .11 and 9 respectively. We calculated the average response and standard deviation over 10,000 presentations for each aspect ratio. Our calculations showed that the detector's output at an aspect ratio of .11 (i.e., when the spatiotemporal correlation was most noisy due to randomly-replotted dots) was well within the standard deviation of the detector's output at an aspect ratio of 9. Thus, our hop size was small enough to avoid low spatiotemporal correlation, and our procedure of randomly-replotting target dots within the target patch did not introduce enough spatiotemporal noise to differentially activate a local motion detector. Finally, any difference between signal strength should have also been evident in Experiment 4, where target and background motions were in orthogonal directions. We suggest, therefore, that differential signal strength or differential signal correlation cannot account for the effects we observe.

The notion that low-level detectors elongated in the direction of motion could account for our data is also not well supported (see, in particular, Experiments 3 and 4). Both psychophysical and physiological data suggest that first stage motion detectors are circular (Daugman, 1985; Anderson & Burr, 1991; Georgeson & Scott-Samuel, 2000). We suggest that the extended spatial and temporal integration extents that we observe are not occurring at the level of first-stage local motion detectors. Of course, other

integration mechanisms built from combinations of elongated and/or circular detectors (or detectors which overlap instead of tile the visual area) could potentially account for the data. Ruling out all possible combinations of detector shape and interaction is clearly beyond the scope of this paper. Rather, our approach has been to test some of the more straightforward explanations given in previous work and build arguments around a few well-defined hypothetical mechanisms.

Our data are consistent with the general notion of cooperativity among low-level motion detectors. As described in the introduction, many have suggested that low-level detectors maintain connections with one another, forming a network in which information is propagated to nearby units with similar directional tuning. Such models accommodate the gross features of our data: they predict a motion integration effect that is directionally-sensitive and dependent on the outputs of first-stage motion detectors. There are some aspects of these models, however, that are inconsistent with our data. The models have spatial “neighborhoods” of interaction between motion detectors, which are typically thought to be of equal extent in all directions (Chang & Julesz, 1984; Gurney & Wright, 1996). Although a circular region of interactions is ideal when the task is motion coherence, our data suggest that interactions occur *in the direction of motion*. In other words, we find that facilitation occurs along a path, not simply over an area. This idea is related to the association field proposed for contour linking across space (Field, Hayes & Hess, 1993). Simply modifying a cooperative motion model so that interactions between the low level detectors are along the direction of motion rather than isotropic fails to address a related issue. The models characterize interactions as cooperative (excitatory) over small spatial extents and competitive (inhibitory) over large spatial extents (Grzywacz, Watamaniuk & McKee,

1995; Lorenceau & Zago, 1999). Even if these interactions occurred in the direction of motion, this notion is hard to reconcile with our data. Assuming proper scaling, this kind of interaction among motion detectors would lead to lowest thresholds as the aspect ratio of our stimulus increases to 1; as the stimulus is elongated in the direction of motion, thresholds would increase. Thus, the idea of cooperative/competitive interactions would predict that discrimination thresholds are a U-shaped function of aspect ratio, and we find no evidence for this. Additionally, the cooperative/competitive arrangement predicts equivalent thresholds for same- and orthogonal- directions of motion in the target and background areas, because the target area in both cases is identical. Either inhibitory interactions play no role in the kind of motion integration observed here, or the spatial scale of competitive interactions is larger than nine degrees (the maximum spatial extent of our stimulus). Thirdly, our data demonstrate that motion integration takes at least 50 ms to develop. We suggest that the integration mechanism cannot be engaged until after low-level motion detectors start to respond, creating the “temporal threshold” present in our data (Figure 5). While our data indicate a lower bound for the time course of the integration mechanism, they do not indicate an upper bound. Other work that has shown very long integration times (on the order of several seconds) for some motion sequences (Todd & Norman, 1995).

Some treatments of motion cooperativity suggest that integration effects similar to those observed here can be explained by facilitation in the direction of motion (Snowden & Braddick ,1989; Yuille & Grzywacz, 1998). The problem with this notion is that it is unclear what kind of signal would be propagated. The work of McKee, Verghese & Yuille (1998) argues against the straightforward idea that a contrast signal is propagated. Recently, Verghese & McKee (2000) have suggested that motion trajectory

effects may be mediated by both a reduction in uncertainty and an increase in sensitivity at subsequent locations along the motion path. The increase in sensitivity is similar to the idea of facilitation along the direction of motion. The reduction in uncertainty suggests that the beginning of a motion sequence predictively constrains the remainder of the sequence, making some motions more likely than others and thereby reducing the physical evidence necessary for detection or discrimination. The mechanism of motion integration, then, is an increase in motion sensitivity along the motion path combined with a reduction in the number of motion detectors that must be monitored. These preliminary results suggest that if stimulus configuration effects are mediated by a higher-level network, the network is best thought of as propagating a measure of the probability of motion in the direction of the initial motion sequence.

FIGURE CAPTIONS

Figure 1. Schematics of mechanisms that could account for integration in the direction of motion. The rectangles represent a fixed image area (the "signal patch") containing dots translating either coherently or at a faster speed than dots outside of the rectangle. The rectangles are oriented either perpendicular (left column) or parallel (right column) to the direction of motion. The circles and ellipses represent low-level motion detectors sensitive to the direction and speed of the translating dots. A) Multiple detector scales. If one assumes that low-level motion detectors are circular, neither signal patch orientation gains an advantage from combining the outputs of multiple detector sizes. B) Elongated detectors. The ellipse represents a low-level detector elongated in the direction of motion. Such a detector would respond differentially depending on the orientation of the signal patch relative to the direction of motion. An elongated detector gets a stronger motion signal when embedded within the signal patch extended in the direction of motion (right column). C) Second-level facilitation in the direction of motion. The lines represent connections between low-level motion detectors tuned to the same direction and speed. When a detector is activated, it spreads activity in the direction of motion. See text for further explanation.

Figure 2. Representation of the stimuli and procedure. Each row represents the two intervals of a single trial. The dark outline, which did not appear in the stimulus display, shows the "signal patch" area. The signal patch appeared equally often in each interval. Dots in the "background" were either randomly-replotted (Experiment 1) or translated at 10 degrees/second (Experiments 2-4). Dots within the rectangular signal patch either moved coherently in some proportion (Experiment 1) or translated more quickly (Experiments 2-4). The direction of motion (rightward or leftward) was chosen randomly for each interval. Top row: Signal patch aspect ratio = .11. Bottom row: Signal patch aspect ratio = 9.

Figure 3. Experiment 1 coherence thresholds plotted as a function of aspect ratio for three observers. Error bars are standard errors of the mean. Aspect ratio (AR)

increases along the abscissa; small values indicate a signal patch oriented perpendicular to the direction of motion, and large values indicate a signal patch oriented parallel to the direction of motion. Thresholds decrease dramatically as the signal patch becomes extended in the direction of motion.

Figure 4. Experiment 2 speed discrimination thresholds plotted as a function of aspect ratio for five observers. Error bars are standard errors of the mean. Thresholds again decrease as the signal patch becomes extended in the direction of motion.

Figure 5. Speed discrimination thresholds and standard errors of the mean for four observers in Experiment 3. Stimulus durations of 100 ms are shown with triangular symbols. Data from Experiment 2, collected at a duration of 200 ms, are re-plotted here with circular symbols. The 400 ms duration data is plotted with square symbols. For clarity, a straight line has been fit to each data series. For 3 of 4 observers, thresholds at 100 ms are unaffected by stimulus configuration.

Figure 6. Experiment 4 speed discrimination thresholds for displays in which signal and background dots moved in orthogonal directions. The data are plotted as before. Thresholds are not as strongly affected by stimulus configuration when signal and background dots move in different directions.

REFERENCES

- Anderson, S. J., & Burr, D. C. (1991). Spatial summation properties of directionally selective mechanisms in human vision. *Journal of the Optical Society of America A*, 8, 1330-1339.
- Anstis, S., & Ramachandran, V. S. (1987). Visual inertial in apparent motion. *Vision Research*, 27(5), 755-764.
- Bravo, M. J., & Watamaniuk, S. N. (1995). Evidence for two speed signals: A coarse local signal for segregation and a precise global signal for discrimination. *Vision Research*, 35(12), 1691-1697.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12), 4745-4765.
- Chang, J. J., & Julesz, B. (1984). Cooperative phenomena in apparent movement perception of random-dot cinematograms. *Vision Research*, 24(12), 1781-1788.
- Daugman, J. G. (1985). Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional cortical filters. *Journal of the Optical Society of America A*, 2, 1160-1169.
- Field, D. J., Hayes, A., & Hess, R. F. Contour integration by the human visual system: evidence for a local “association field”. *Vision Research*, 33(2), 173-193.
- Fredericksen, R. E., Verstraten, F. A. J., & van de Grind, W. A. (1994). Spatial summation and its interaction with the temporal integration mechanism in human motion perception. *Vision Research*, 34(23), 3171-3188.

- Fredericksen, R. E., Verstraten, F. A. J., & van de Grind, W. A. (1997). Pitfalls in estimating motion detector receptive field geometry. *Vision Research*, 37(1), 99-119.
- Georgeson, M. A., & Scott-Samuel, N. E. (2000). Spatial resolution and receptive field height of motion sensors in human vision. *Vision Research*, 40(7), 745-758.
- Grzywacz, N. M., Watamaniuk, S. N. J., & McKee, S.P. (1995). Temporal coherence theory for the detection and measurement of visual motion. *Vision Research*, 35(22), 3183-3203.
- Gurney, K. N., & Wright, M. J. (1996). A model for the spatial integration and differentiation of velocity signals. *Vision Research*, 36(18), 2939-2995.
- Lappin, J. S., & Bell, H. H. (1976). The detection of coherence in moving random-dot patterns. *Vision Research*, 16, 161-168.
- Lorenceau, J., & Zago, L. (1999). Cooperative and competitive spatial interaction in motion integration. *Visual Neuroscience*, 16, 755-770.
- McKee, S. P. (1981). A local mechanism for differential velocity detection. *Vision Research*, 21, 491-500.
- McKee, S. P., Verghese, P., & Yuille, A. L. (1998). Trajectory facilitation is not mediated by contrast changes. *Investigative Ophthalmology & Visual Science*, 39, S228.
- Newsome, W. T., & Pare, E. B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience*, 2201-2211.
- Snowden, R. J., & Braddick, O. J. (1989). The combination of motion signals over time. *Vision Research*, 29(11), 1621-1630.

- Snowden, R. J., & Braddick, O. J. (1991). The temporal integration and resolution of velocity signals. *Vision Research*, 31(5), 907-914.
- Todd, J. T., & Norman, J. F. (1995). The effects of spatiotemporal integration on maximum displacement thresholds for the detection of coherent motion. *Vision Research*, 35(16), 2287-2302.
- van Doorn, A. J., & Koenderink, J. J. (1984). Spatiotemporal integration in the detection of coherent motion. *Vision Research*, 24(1), 47-53.
- Vergheze, P., Watamaniuk, S. N. J., McKee, S. P., & Grzywacz, N. M. (1999). Local motion detectors cannot account for the detectability of an extended trajectory in noise. *Vision Research*, 39, 19-30.
- Vergheze, P., & McKee, S. P. (2000). Reduced uncertainty accounts for the enhanced sensitivity to motion trajectories. *Abstr. Soc. Neurosci.*, 687.6.
- Vergheze, P., McKee, S. P., & Grzywacz, N. M. (2000). Stimulus configuration determines the detectability of motion signals in noise. *Journal of the Optical Society of America A*, 17(9), 1525-1534.
- Watamaniuk, S. N. J., & Duchon, A. (1992). The human visual system averages speed information. *Vision Research*, 32(5), 931-941.
- Watamaniuk, S. N. J., McKee, S. P., & Grzywacz, N. M. (1995). Detecting a trajectory embedded in random-direction motion noise. *Vision Research*, 35(1), 65-77.
- Watson, A. B. (1979). Probability summation over time. *Vision Research*, 19, 515-522.
- Watson, A. B., & Ahumada Jr., A. J. (1985). Model of human visual-motion sensing. *Journal Opt. Soc. Am. A*, 2(2), 322-341.

Watson, A. B., & Turano, K. The optimal motion stimulus. *Vision Research*, 35(3), 325-336.

Williams, D., Philips, G., & Sekuler, R. (1986). Hysteresis in the perception of motion direction as evidence for neural cooperativity. *Nature*, 324, 253-255.

Yuille, A. L. & Grzywacz, N. M. (1998). A theoretical framework for visual motion. In T. Watanabe (Ed.) *High-level Motion Processing: From Computational, Neurophysiological, and Psychophysical Perspectives*. MIT Press, Cambridge.

Figure 1.




Figure 2.

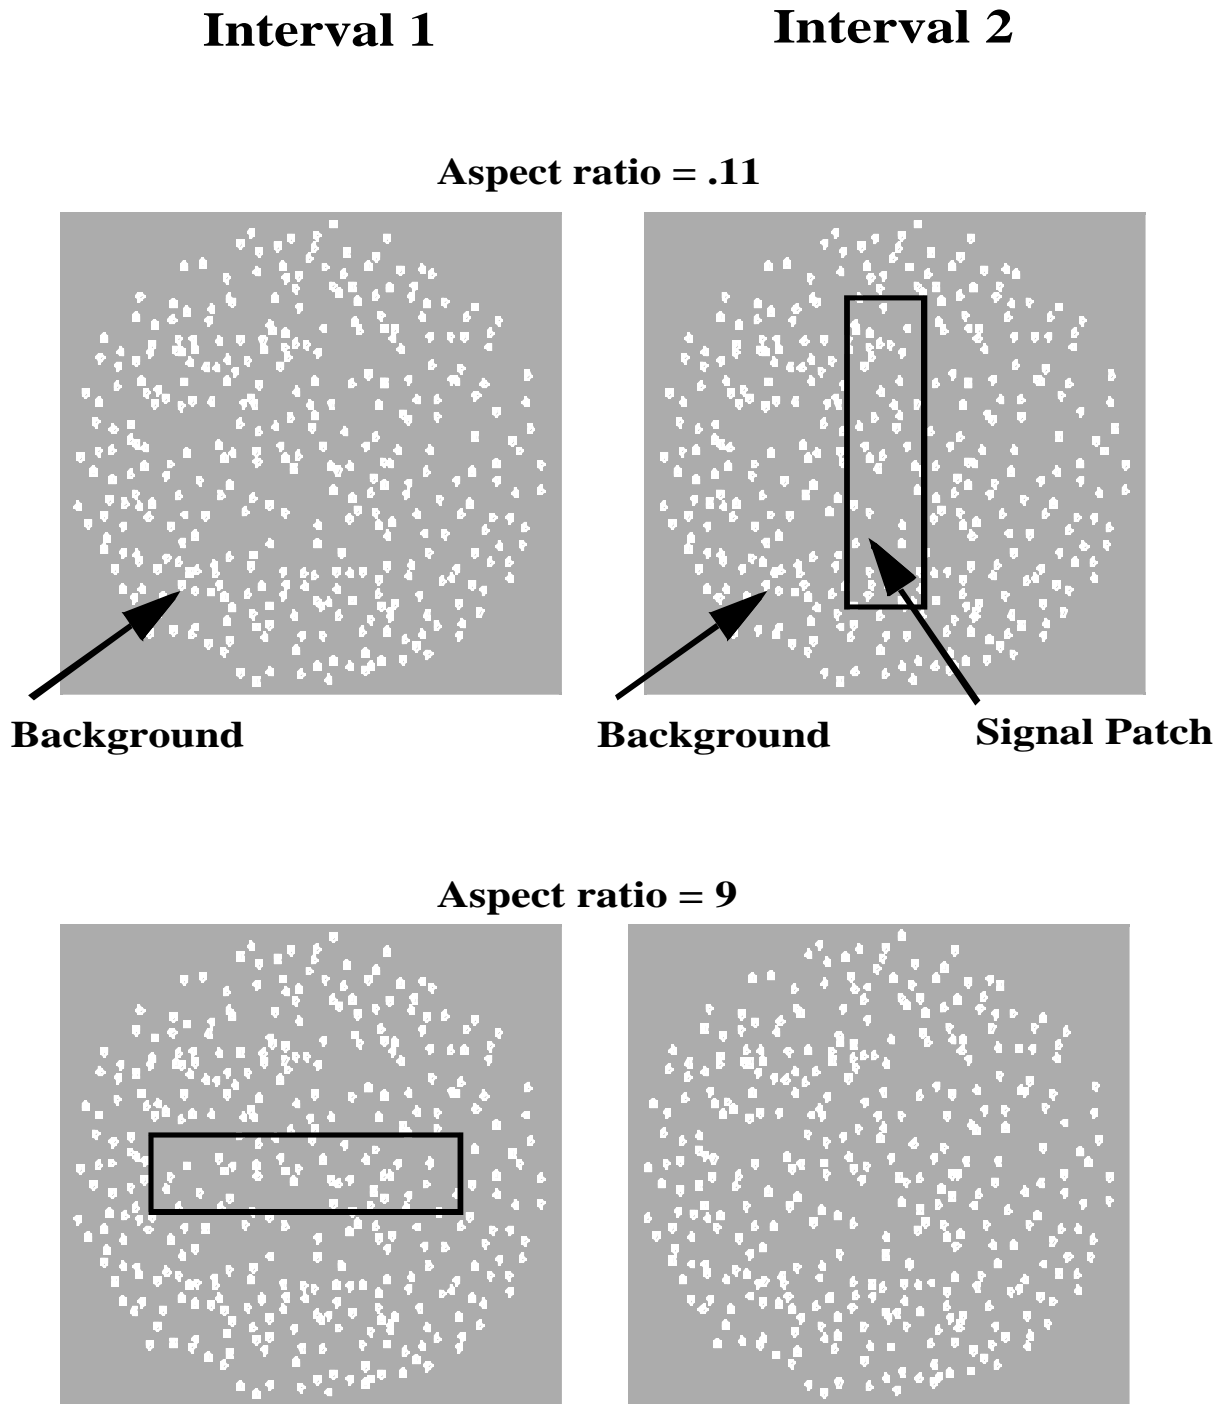


Figure 3.




Figure 4.

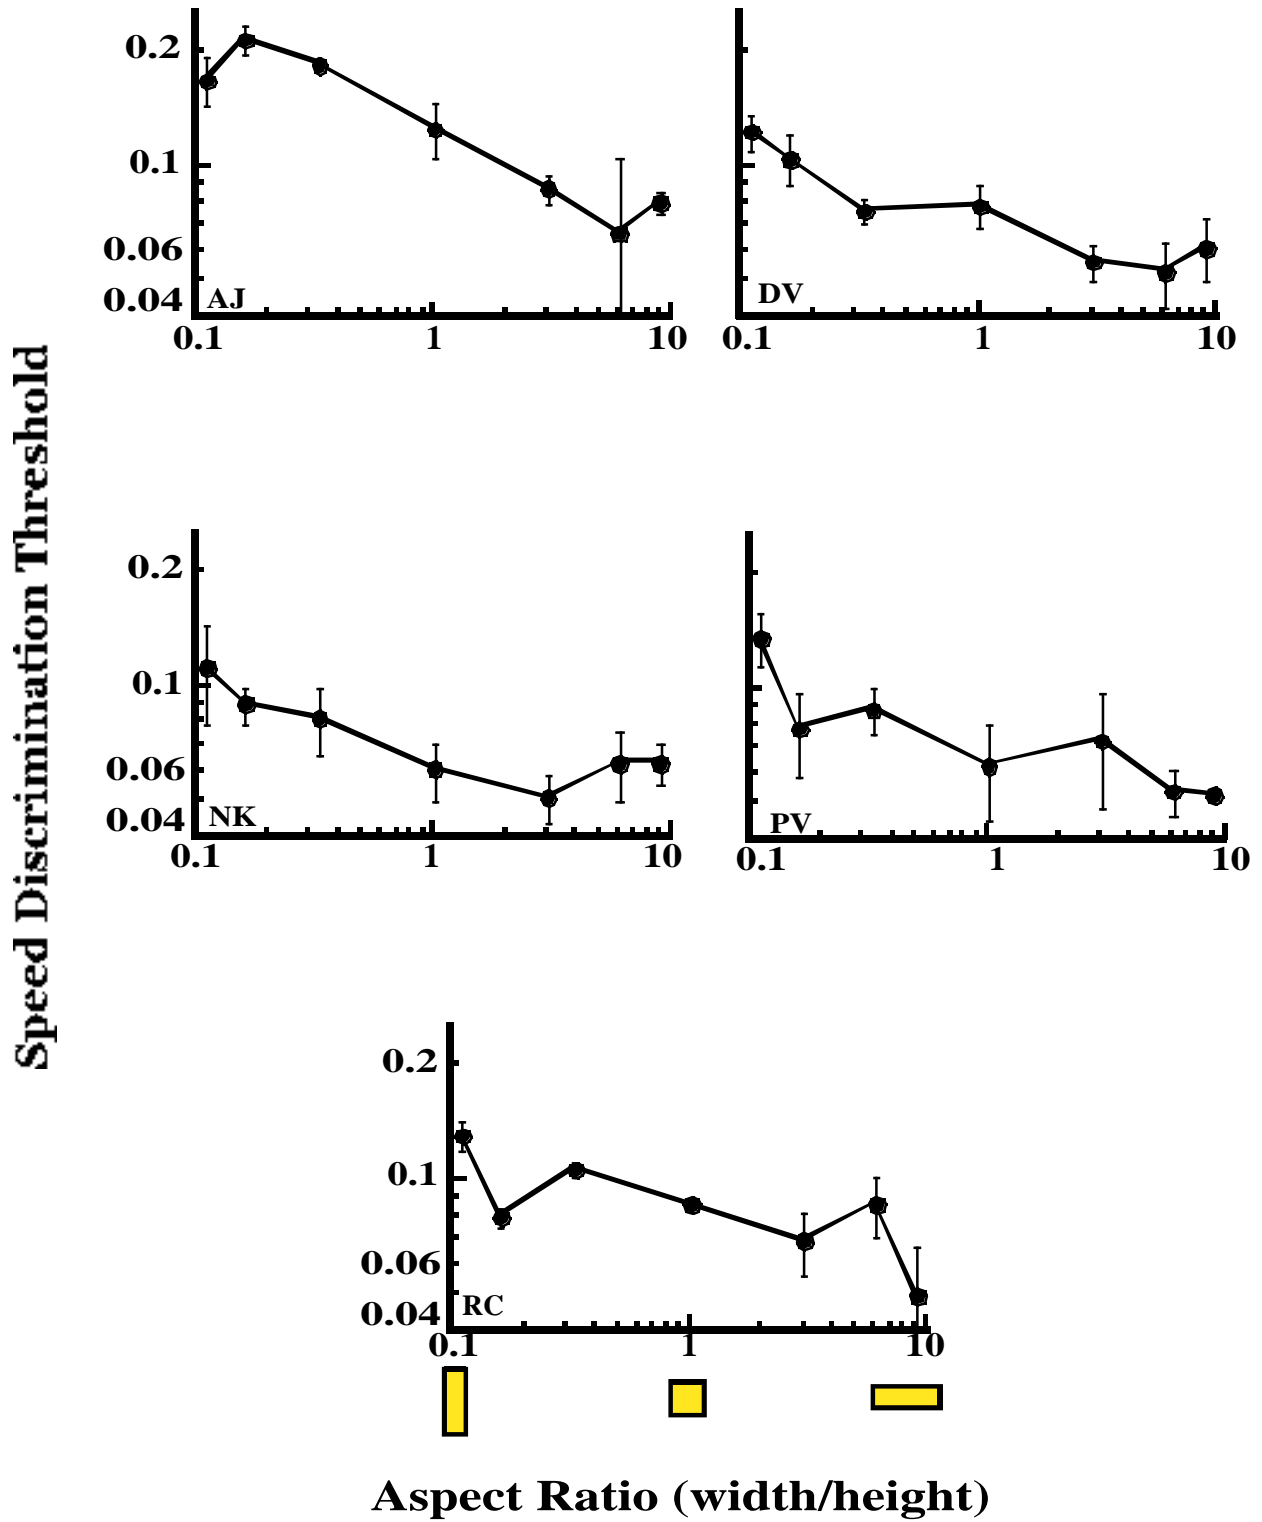


Figure 5.

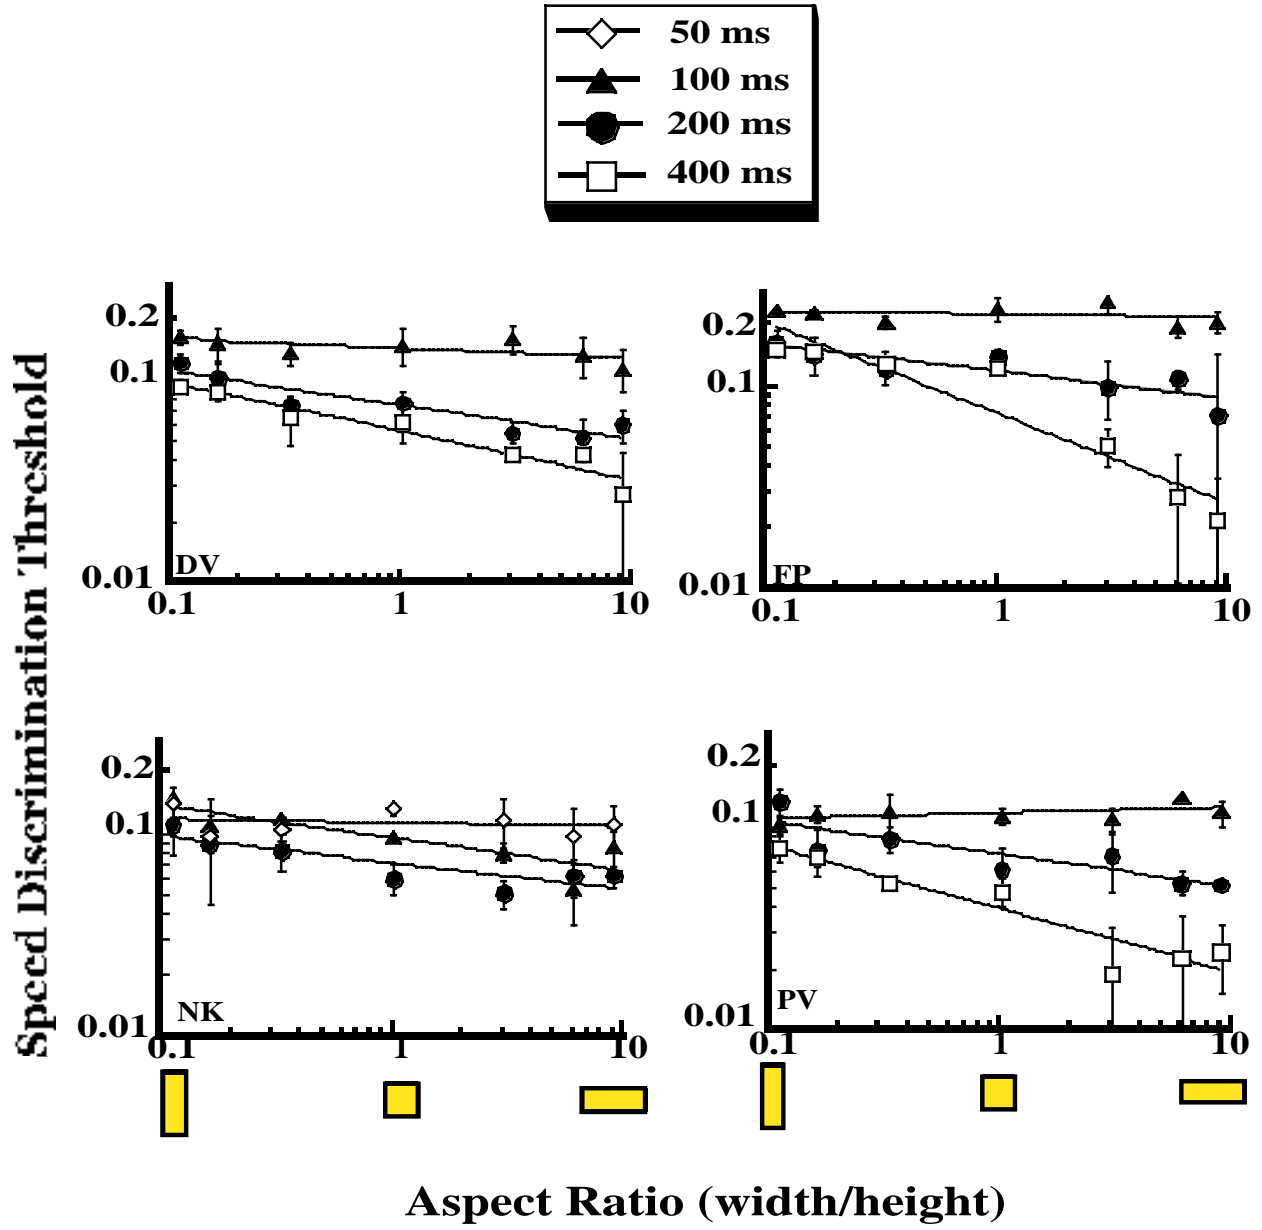


Figure 6.

