Pergamon

0042-6989(94)00266-5

Vision Res. Vol. 35, No. 12, pp. 1691–1697, 1995
Copyright © 1995 Elsevier Science Ltd
Printed in Great Britain. All rights reserved
0042-6989/95 $9.50 + 0.00

Evidence for Two Speed Signals: a Coarse Local Signal for Segregation and a Precise Global Signal for Discrimination

MARY J. BRAVO,*† SCOTT N. J. WATAMANIUK*

Received 8 December 1993; in revised form 1 October 1994

In calculating the precise speed of an object, the visual system must integrate motion measurements across time and space while keeping motion measurements from different objects separate. We examined whether an initial coarse estimate of local speed may be used to segregate the motions of different objects prior to a precise calculation of object speed. Our stimuli consisted of 256 dots that moved upward at two speeds. In Expt 1, each dot alternated between the two speeds every 133 msec. When the speed alternations were asynchronous across dots, subjects saw two transparent surfaces moving at different speeds and their ability to discriminate changes in the slow speed were unaffected by the presence of the fast speed. This experiment suggests that before integration, motion measurements may be segregated according to speed. We sought more conclusive evidence for this claim in Expts 2 and 3. In Expt 2, dots with 33 msec lifetimes were used to generate the two speeds. Although individual dots permitted only crude speed discrimination, subjects perceived this stimulus as two surfaces moving at different speeds and they precisely judged the slower speed. Apparently, the coarse local signals generated by the slow dots were segregated from those of the fast dots and then separately integrated to produce a precise speed signal. In Expt 3, the dots again moved at two speeds, but each speed was generated by a range of spatial and temporal displacements. Once more, subjects saw two surfaces and precisely judged the speed of the slower surface, demonstrating that segregation may be based solely on differences in local speed. We conclude that the visual system calculates two speed signals, one speed signal is coarse, local and used for segregation and the second signal is precise, global and used for speed discrimination.

Human psychophysics Local velocity Speed discrimination Motion transparency

INTRODUCTION

Human observers can judge an object's velocity with remarkable precision: under optimal conditions observers can detect a 5% difference in the speed, and a 1 deg difference in the direction of two moving objects (McKee, 1981; McKee, Silverman & Nakayama, 1986; Levinson & Sekuler, 1976; Watamaniuk & Duchon, 1992; Watamaniuk & Sekuler, 1992). Although much research has focused on this remarkable aspect of visual function, it is still unclear when and how velocity is calculated in motion processing (Braddick, 1993).

In current models of human motion processing, the earliest motion-selective units are not selective for velocity (Adelson & Bergen, 1985; van Santen & Sperling, 1984; Watson & Ahumada, 1985). Instead these motion detectors, like those in the primary visual cortex of the cat, are selective for spatial frequency (SF), temporal frequency (TF), orientation, and direction (Tolhurst & Movshon, 1975; Holub & Morton-Gibson, 1981). Because motion detectors are tuned to TF and SF independently, and not to a TF:SF ratio, they are not tuned to speed. In addition, because early motion detectors are orientation selective, they respond mainly to the component of the stimulus motion that is orthogonal to their preferred orientation.

Although individual detectors are not selective for velocity, a local velocity may be calculated by comparing the activities of motion-selective units with different TF and SF tunings but common receptive fields. Several computational models that calculate a local velocity have been proposed (Heeger, 1988; Grzywacz & Yuille, 1990; Heeger & Simoncelli, 1992). Constraining the velocity calculation to use only local information reduces the chance that the motions of different objects will be combined in the velocity estimate. Nonetheless, it may be impossible to calculate a precise, unambiguous velocity using only local information. To calculate an

---

*The Smith-Kettlewell Eye Research Institute, 2232 Webster Street, San Francisco, CA 94115, U.S.A.

†To whom all correspondence should be addressed at present address: University of Pennsylvania, Computer and Information Sciences, GRASP Laboratory, 3401 Walnut Street #334C, Philadelphia, PA 19104, U.S.A.

unambiguous velocity, it is necessary to measure motion in at least two directions. If, locally, an object's contours have a single orientation, then the local velocity measurement reflects only the motion in the direction orthogonal to that contour (Fennema & Thompson, 1979; Marr, 1982; Hildreth, 1984). Even when local velocity is unambiguous, it may still be imprecise. The precision of a velocity estimate depends in part on the density with which the motion-selective units sample the range of visible SFs and TFs (Grzywacz & Yuille, 1990). If, locally, this sampling density is low, then a precise velocity estimate may require integrating motion measurements from different spatial locations. The precision of a velocity estimate also depends on the degree to which noise is correlated across motion detectors. If, locally, this noise is highly correlated, then a precise velocity estimate will require averaging measurements across space and time.

Empirical evidence against a local velocity signal comes from experiments demonstrating the visual system's insensitivity to acceleration. Although subjects can detect a 5% difference in the speed of two stimuli separated by a temporal interval, they cannot detect an acceleration that produces a similar change in the speed of a single stimulus (Gottsdanker, 1965; Snowden & Braddick, 1991; Werkhoven, Snippe & Toet, 1992). In addition, speed thresholds are elevated when the target motion is a segment of an accelerating trajectory compared to when the target motion is presented in isolation (Bowne, McKee & Glaser, 1989). To account for such results, Snowden and Braddick (1991) and Bowne *et al.* (1989) have posited mechanisms that integrate motion information over space and time prior to the calculation of velocity.

Snowden and Braddick (1991) suggest that the visual system's sluggish response to changes in velocity is due to interactions between motion detectors. These putative interactions facilitate the detection of similar speeds and inhibit the detection of dissimilar speeds. Because these interactions have a rapid onset time and a slow decay time, one velocity will suppress a second velocity that is presented either simultaneously or after a short delay.

Alternatively, Bowne *et al.* (1989) propose a mechanism designed to integrate motion along object trajectories. They postulate units sensitive to temporal delays in the activity of spatially offset motion detectors. These units act as a second layer of motion detectors that respond, for example, to a left-to-right sequence in the activation of first layer detectors that are themselves tuned to left-to-right motion. Thus, Bowne *et al.* suggest that integration occurs whenever motion detectors are stimulated in the appropriate spatial and temporal sequence.

Note that in both accounts the spatial and temporal integration of motion measurements occurs before velocity is calculated. This integration across motion detectors prohibits the calculation of a local velocity. Thus, in explaining their empirical results, these authors implicitly reject the existence of a local velocity. But does the insensitivity of the visual system to acceleration necessarily rule out the existence of a local velocity signal? Werkhoven *et al.* (1992) suggest that the visual system does measure local velocity, and they explain our insensitivity to acceleration in terms of a mechanism that measures the variance of these local velocities over time. We offer a different account.

We interpret the insensitivity of the visual system to acceleration in terms of two velocity calculations. The first calculation produces a coarse, and sometimes ambiguous, local signal. These local signals are then used to segregate the motions of different objects. Motions assigned to the same object are integrated across space and time to produce a precise and unambiguous estimate of the object's velocity. When the visual system is presented with a rigid pattern that changes velocity over time, it calculates the local velocity of the pattern over limited spatial and temporal intervals. But because the stimulus is interpreted as a single accelerating object, the local velocities are then integrated across space and time. This integration causes the subject to be relatively less sensitive to changes in the stimulus velocity. This account may be tested by using stimuli that the visual system interprets as multiple objects.

EXPERIMENT 1

An experiment by Snowden and Braddick (1991) provides the best introduction to our first experiment. Snowden and Braddick presented their subjects with two stimuli: a dot pattern that moved at a constant speed and a dot pattern that alternated between two speeds at a frequency of 2 Hz or greater. To discriminate between these stimuli, subjects required at least a 30% difference between the two speeds in the alternating stimulus. Snowden and Braddick contrast this 30% threshold with the 5–6% threshold found when subjects discriminated differences in the speed of two random-dot patterns each presented for 100 msec with a 500 msec ISI. It was this five-fold difference in the speed discrimination thresholds that led Snowden and Braddick to postulate slowly decaying interactions between motion detectors.

In our first experiment, we also used a stimulus composed of random dots that alternated between two speeds. But instead of having all dots change speed simultaneously, we asynchronized these changes. As we describe below, this slight modification produces a radically different result.

Stimulus

Our stimulus was a random-dot cinematogram displayed on a Tektronix oscilloscope (P4 phosphor). On the first frame of the cinematogram the positions of the 256 dots were chosen randomly; on successive frames all dots were displaced vertically. When a dot reached the top of the display it wrapped around to the bottom of the display. Each dot moved along a continuous vertical trajectory, alternating every 133 msec between a slow speed (either 5.4, 5.7, 6.0, 6.3, or 6.6 deg/sec) and a fast speed (either 18.9, 20.0, 21.0, 22.0 or 23.1 deg/sec).

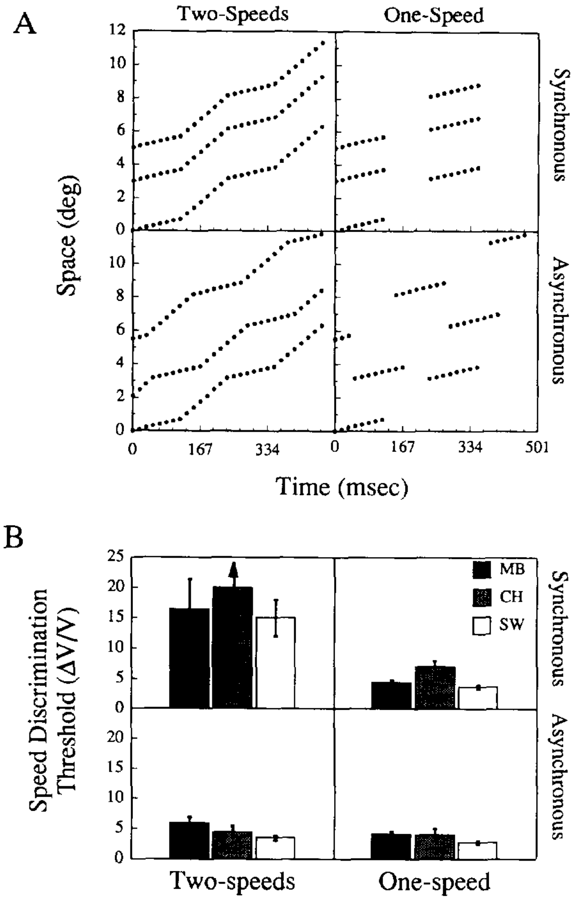

FIGURE 1. (A) Plot of the dot motions in the four stimulus conditions of Expt 1. Actual two-speed and one-speed stimuli contained 256 and 128 dots respectively. Note that the motions of individual dots are identical in the synchronous and asynchronous conditions. (B) Speed discrimination thresholds for three subjects on the four stimulus conditions depicted in (A).

These instantaneous speed alternations were either synchronous or asynchronous across dots.

Synchronous alternations. In the synchronous alternation condition, all dots moved at one speed for 133 msec, then all dots switched to the other speed for 133 msec, then all dots switched to the first speed for 133 msec, and so on [Fig. 1(A), upper left]. Thus, although both speeds were presented during a trial, only one speed was presented at any given moment. This display is analogous to that of Snowden and Braddick (1991). We also created "one-speed" displays in which the fast speed was eliminated. In these synchronous one-speed displays all 256 dots moved at the slow speed for 133 msec, disappeared for 133 msec and then reappeared moving at the slow speed for 133 msec [Fig. 1(A), upper right]. During the blank 133 msec interval, each dot was displaced by the distance the dot would have traveled had it continued to move for 133 msec at the fast speed.

Asynchronous alternations. In the asynchronous alternation condition, the behavior of the individual dots was the same as in the synchronous condition; each dot alternated between the two speeds every 133 msec. But now the speed alternations were no longer simultaneous across dots. Instead, half of the dots started with the slow speed, half started with the fast speed, and some dots changed speed on every frame [Fig. 1(A), lower left]. Thus, in the asynchronous condition both speeds were always present and each had a spatial distribution that was random and constantly changing. Asynchronous one-speed stimuli were also created. The behavior of individual dots was the same as in the synchronous one-speed stimulus: each dot disappeared and then reappeared every 133 msec, but these disappearances and reappearances were no longer simultaneous across dots [Fig. 1(A), lower right].

For all conditions, the stimulus duration varied randomly between 450 and 533 msec. Subjects viewed the displays binocularly from a distance of 57 cm. The display subtended 10 × 10 deg, and the dots subtended 4.2 min arc. The screen luminance was 33 cd/m^2 and the dots had a space-averaged luminance of 68 cd/m^2.* The stimulus frame rate was 60 Hz.

Procedure

Speed discrimination thresholds were measured using the single stimulus variant of the method of constant stimuli (McKee, 1981). While fixating a central spot, subjects initiated a trial by pressing a button. After the stimulus was presented, the subject pressed one of two buttons to signal whether the slow speed was greater or less than the mean slow speed across trials. Feedback was given. Although only one slow and one fast speed were presented during each trial, all possible pairings of the five slow and five fast speeds were presented during a block of trials. Since each slow speed was randomly paired with five fast speeds the relative motion of the fast and slow dots did not provide a reliable cue to the magnitude of the slow speed. A block of trials consisted of 20 practice trials followed by 250 experimental trials: 10 trials for each pairing of fast and slow speeds, presented in random order. Subjects repeated each condition until their performance plateaued and speed discrimination thresholds were based on a final block of trials. After running each condition, subjects described what they saw.

For each slow speed, the number of fast responses was used to generate a psychometric function. The data were fit with a cumulative normal distribution using Probit analysis. Speed discrimination thresholds were defined as half the speed difference necessary to change performance levels from 25% to 75%. Standard errors were derived from the variance of the psychometric function.

Subjects

The two authors and a paid, naive subject participated in these experiments. All three subjects had participated in numerous psychophysical experiments and had normal, or corrected-to-normal, vision.

---

*This space-averaged luminance was measured using a matrix of dots with a center-to-center spacing of 4.8 min arc and a frame rate of 60 Hz.

Results and discussion

When the stimulus contained only one speed, speed discrimination performance was similar for the synchronized and asynchronized conditions. In both cases speed discrimination thresholds were low, ranging from 3% to 7% [Fig. 1(B), right]. In contrast, when the stimulus contained two speeds, the two conditions gave very different results.

When the dots alternated synchronously between the two speeds, speed discrimination thresholds for the slower speed were greatly affected by the presence of the fast speed. Because one subject was unable to discriminate the largest speed change we used, 20%, we were unable to measure her threshold. The thresholds for the other two subjects were 16.4% and 15%. When the dots alternated asynchronously between the two speeds, speed discrimination thresholds for the slower speed returned to the levels of the one-speed condition. The discrimination thresholds for the asynchronous two speed condition ranged from 3% to 6%.

The subjective reports of the subjects mirrored their speed discrimination performance. In the synchronous two-speed condition, subjects reported seeing a single surface lurching across the screen, and the speed of this surface appeared to change in a graded way between two values. In the asynchronous condition, subjects reported seeing a distinct surface moving at the slow speed and a second surface moving at the fast speed.* Even though each dot moved on a continuous trajectory, the surfaces appeared to twinkle, presumably because as a dot changed speed it disappeared from one surface and reappeared on the other surface.

In sum, the synchronous condition replicates earlier studies, but the asynchronous condition refutes the conclusion that has been drawn from these studies. The asynchronous condition clearly shows that motion measurements are not automatically integrated along trajectories, nor are they automatically integrated within a spatial and temporal window. Instead, the initial motion measurements are used to calculate a local speed and integration occurs after these local speeds have been used to segregate the stimulus.

EXPERIMENT 2: INTEGRATION AFTER SEGREGATION?

The previous experiment demonstrated that two very different speeds may be segregated and the speeds judged independently. However, this experiment did not test whether the slow speeds that had been segregated from different dot trajectories were integrated for a precise speed estimate. Each dot maintained a constant speed for 133 msec, a duration sufficient for good speed discrimination (McKee, 1981). Subjects could have segmented the trajectories of the dots according to speed and then based their speed judgments on a single 133 msec segment. In this experiment we used dots with durations so brief that they were poor stimuli for speed discrimination. This allowed us to determine whether, following segregation, the visual system integrates motion signals to arrive at a precise speed measurement.

Methods

The stimulus was similar to the asynchronous alternation condition of the previous experiment except that the dots no longer moved across the screen with continuous trajectories. Instead, when a dot changed speed it was replotted in a random location. Thus a dot moved upward at the slow speed for 133 msec, was displaced to a new location, and then moved upward at the fast speed for 133 msec. Each of the 256 dots started at a random point in this cycle. In other words, the stimulus was composed of dots with 133 msec lifetimes. We also generated stimuli composed of dots with 66 and 33 msec lifetimes respectively. [Figure 2(A) corresponds to a 33 msec stimulus.] For each dot lifetime, a one-speed stimulus was generated by eliminating the 128 dots moving at the fast speed.

To allow us to determine whether speed discrimination of these multiple-dot stimuli was based on a single dot, stimuli consisting of a single dot were generated. The dot moved upward for 133, 66 or 33 msec at one of the five slow speeds. The starting position of the dot was 0.5 deg to the left and 1 ± 0.33 deg below the fixation mark. The random starting position of the dot prevented subjects from basing their speed judgments on the distance the dot traveled relative to the fixation spot.

The procedure for this experiment was the same as in the previous experiment. On each trial subjects judged whether the slow speed was greater or less than the average slow speed across trials. Each condition was run in a separate block of trials.

Results and discussion

The open circles with dashed lines in Fig. 2(B) show the results for the one-dot stimuli. Speed discrimination thresholds improve several-fold as dot lifetime increases from 33 to 133 msec, in replication of McKee (1981), Orban, Wolf and Maes (1984) and Snowden and Braddick (1991). The solid circles of Fig. 2(B) show the results for the stimuli with multiple dots moving at one speed. Speed discrimination thresholds were low even for the shortest dot lifetime, indicating that speed judgments were not based on a single dot. Apparently, the coarse signals from several dots were integrated to arrive at a precise speed estimate. The solid squares of Fig. 2(B) show the results for the stimuli with multiple dots moving at one of two speeds. Again, speed discrimination thresholds were low for all dot lifetimes. Subjects also reported that these stimuli appeared as two surfaces moving at different speeds. We conclude then that coarse local signals may be segregated according to speed before being integrated for a precise speed estimate.

---

*Upon viewing these displays, several individuals reported seeing two global motions and two groups of dots, but not two surfaces. These individuals also report that one-speed displays do not appear as surfaces.

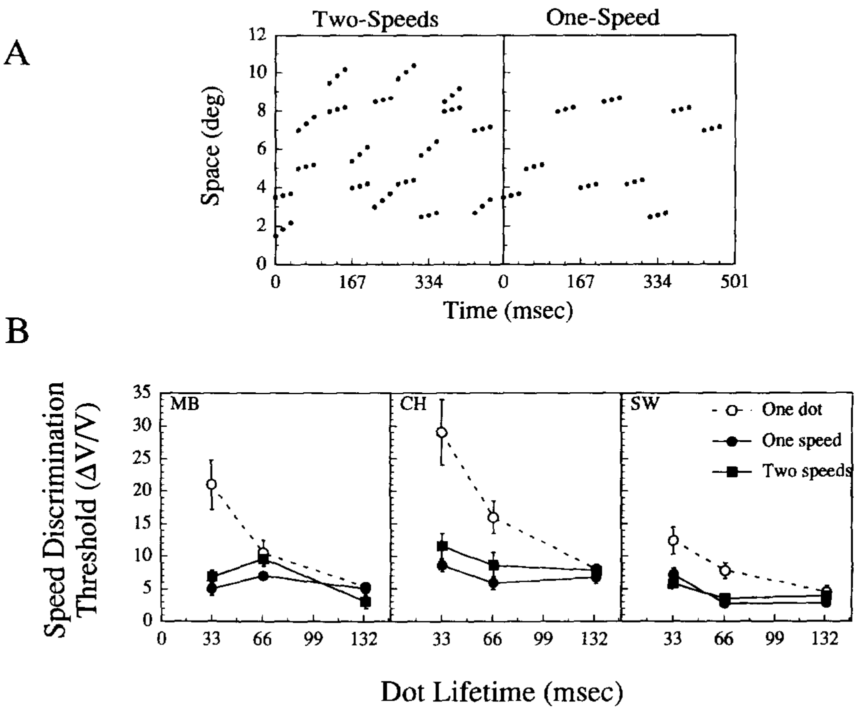

FIGURE 2. (A) Plot of the dot motions in the one-speed and two-speed conditions of Expt 2. (B) Speed discrimination thresholds for three subjects. The one-speed and two-speed conditions are represented in (A), the one-dot condition consisted of a single dot moving at a slow speed.

EXPERIMENT 3: IS SEGREGATION BASED ON SPEED?

In interpreting the two previous experiments, we assumed that segregation was based on speed. However, it is possible that segregation was based on another cue, namely step size (the size of the spatial displacement of the dots from frame to frame). Dots moving at 6 deg/sec had a step size of 0.1 deg/frame, while dots moving at 21 deg/sec had a step size of 0.35 deg/frame. A unit not selective for motion but having a receptive field diameter greater than the step size of the slow dots but less than the step size of the fast dots would generate a greater response to the slow dots. This could cause, for example, slow dots to appear brighter than fast dots and segregation could then be based on this brightness difference. We eliminated cues associated with step size in this final experiment.

Methods

These displays were identical to the asynchronous-alternation condition in Expt 1 except for two changes. First, each dot took four steps at each speed, but the steps varied in their spatial and temporal extent as shown in Fig. 3(A). Second, the faster speeds were reduced to 10.8, 11.4, 12.0, 12.6, or 13.2 deg/sec so that the step sizes for the slow and fast speeds would overlap. Preliminary experiments demonstrated that when step size was constant, subjects could segregate these fast speeds from the slow speeds.

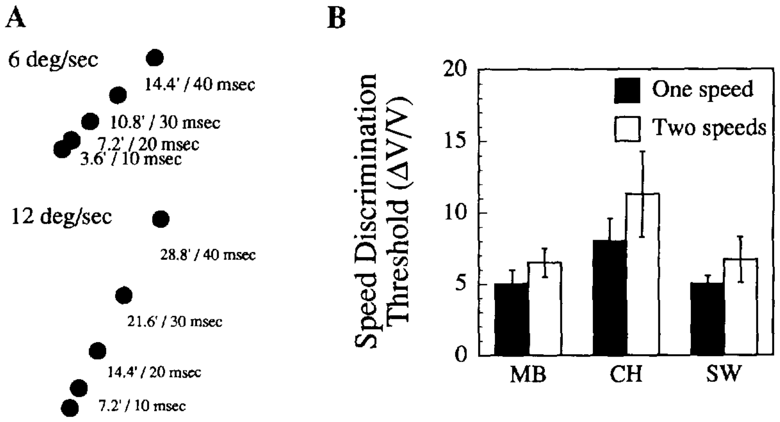

FIGURE 3. (A) Four space-time intervals were used to generate the slow and fast speeds of Expt 3. (B) Speed discrimination thresholds for three subjects on the one-speed and two-speed stimuli where each speed was generated by a range of space-time intervals.

The procedure was the same as in the two previous experiments.

Results and discussion

Figure 3(B) shows the speed discrimination thresholds for the one-speed stimulus (solid bars) and the two-speed stimulus (open bars) in which each speed was generated by four different step sizes. For each subject the speed discrimination threshold for the two-speed stimulus was similar to the one-speed stimulus. Further, the subjects reported that the two-speed stimulus appeared as two surfaces. Because segregation occurred in the absence of the step size cue, we conclude that differences in local speed are sufficient for the segregation of motion measurements.

GENERAL DISCUSSION

These experiments suggest that the visual system resolves the trade-off between the spatial resolution and the precision of the velocity measurement by calculating velocity twice.* The first calculation involves only local motion measurements and produces a coarse signal that is used for segregating the motions of different objects. The second calculation integrates the motion measurements that have been assigned to the same object to produce a precise signal that is used for judgments of object speed.

Several models of the local velocity calculation have been proposed (Heeger, 1988; Grzywacz & Yuille, 1990; Heeger & Simoncelli, 1992) and recently these models have been modified to permit the calculation of two velocities in the same location (Smith & Grzywacz, 1995; Darrell & Simoncelli, 1993). This modification is designed to account for the phenomenon of motion transparency which is often demonstrated by superimposing two grating patterns moving in different directions (Adelson & Movshon, 1982). Under some conditions, observers perceive such stimuli as two transparent surfaces with different motions. Although the dot stimuli used here produce motion transparency, we must note that they do not necessarily require that the visual system calculate two velocities at one location. The density of the dots in these experiments was low (2.6 dots/deg^2 on average). Thus in many locations the two velocities may not have overlapped, and these locations may have provided sufficient information for the segregation of the stimulus into two objects.

The simplest strategy for segregating these displays is to generate a velocity histogram by collecting local velocities across space. Segregation could then be based on clusters of activity in this histogram. Experiment 1 indicates that in generating this histogram, the visual system does not collect local velocities over time. Recall that in this experiment, segregation did not occur in the synchronous alternation condition in which the two velocities were presented sequentially.

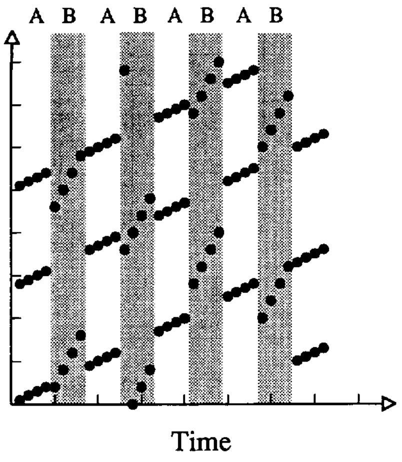

FIGURE 4. Schematic representation of the van Doorn and Koenderink (1982) stimulus which consisted of two drifting random-dot patterns which were alternately presented. The actual patterns had 50% black and 50% white dots.

At first glance, the conclusion that segregation is not based on comparing local velocities over time appears to be contradicted by a 1982 study by van Doorn and Koenderink. They found that a stimulus that alternated synchronously between two speeds appeared to segregate. However their stimulus differed from ours in a critical respect. They used two cinematograms, each a random-dot pattern (50% black and 50% white dots) drifting rightward. The two cinematograms were interlaced: the first *N* frames of cinematogram A were shown, then the second *N* frames of cinematogram B, then the third *N* frames of A etc. (Fig. 4). Note that this stimulus differs from our synchronous stimulus which contained a single dot pattern that alternated between two speeds. van Doorn and Koenderink report that their subjects saw two transparent surfaces when the alternations between the two cinematograms occurred at least every 40 msec. Since motion detectors are assumed to have an integration time greater than 40 msec, these units would detect motion across successive presentations of each cinematogram, and as a result, both speeds would be visible simultaneously. Thus speed segregation may require that both speeds be visible simultaneously, where simultaneous is defined by the integration time of motion detectors.

While the simple segregation scheme described above could readily segregate two transparent translating surfaces, it would fail to segregate many other stimuli. Most naturally occurring surfaces produce a velocity distribution that is more complex than that of the stimuli used here. Surfaces that rotate, dilate or are slanted with respect to the observer produce a range of velocities which may include velocities common to other surfaces in the image. If the visual system can segregate such surfaces then it must use a different strategy.

Once local velocities are segregated, the local velocities that have been assigned to the same object are integrated

---

*In agreement with computational models of motion processing, we assume that speed is not calculated independently of direction. So although these experiments examine segregation based on local speed, in this discussion we consider segregation based on local velocity.

across time and space to produce a precise estimate of object velocity. Numerous computational models of motion processing have explained the need for integrating motion signals, but the models assign this integration to various stages in motion processing (see e.g. Yuille & Grzywacz, 1988; Hildreth, 1984). The experiments reported here indicate that this integration occurs after a local speed calculation. Integrating motion information over space and time allows the visual system to judge velocity with greater precision than would be possible with a single measure of local velocity. We saw this in Expt 2, in which the speed discrimination threshold for a field of dots was significantly better than the threshold for a single dot. But integrating motion information over space and time may also smooth over small changes in local velocity over time producing the insensitivity to acceleration observed by Gottsdanker (1965), Snowden and Braddick (1991) and Werkhoven *et al.* (1992).

REFERENCES

Adelson, E. H. & Bergen, J. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, *2*, 284–299.

Adelson, E. H. & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature*, *300*, 523–525.

Bowne, S. F., McKee, S. P. & Glaser, D. A. (1989). Motion interference in speed discrimination. *Journal of the Optical Society of America A*, *6*, 1112–1121.

Braddick, O. (1993). Segmentation versus integration in visual motion processing. *Trends in Neuroscience*, *16*, 263–268.

Darrell, T. & Simoncelli, E. (1993). Separation of transparent motion into layers using velocity-tuned mechanisms. M.I.T. Media Laboratory Technical Report No. 244.

van Doorn, A. J. & Koenderink, J. J. (1982). Temporal properties of the visual detectability of moving spatial white noise. *Experimental Brain Research*, *45*, 179–188.

Fennema, C. L. & Thompson, W. B. (1979). Velocity determination in scenes containing several moving images. *Computer Graphics and Image Processing*, *9*, 301–315.

Gottsdanker, R. M. (1965). The ability of human operators to detect acceleration of target motion. *Psychological Bulletin*, *53*, 477–487.

Grzywacz, N. M. & Yuille, A. L. (1990). A model for the estimate of local image velocity by cells in the visual cortex. *Proceedings of the Royal Society of London B*, *239*, 129–161.

Heeger, D. J. (1988). Optical flow using spatiotemporal filters. *International Journal of Computer Vision*, *1*, 279–302.

Heeger, D. J. & Simoncelli, E. P. (1992). Model of visual motion sensing. In Harris, L. & Jenkin, M. (Eds), *Spatial vision in humans and robots*. Cambridge: Cambridge University Press.

Hildreth, E. C. (1984). *The measurement of visual motion*. Cambridge, Mass.: MIT Press.

Holub, R. A. & Morton-Gibson, M. (1981). Response of visual cortical neurons of the cat to moving sinusoidal gratings: Response–contrast functions and spatiotemporal integration. *Journal of Neurophysiology*, *46*, 1244–1259.

Levinson, E. & Sekuler, R. (1976). Adaptation alters perceived direction of motion. *Vision Research*, *16*, 779–781.

Marr, D. (1982). *Vision*. San Francisco, Calif.: Freeman.

McKee, S. P. (1981). A local mechanism for differential velocity detection. *Vision Research*, *21*, 491–500.

McKee, S. P. & Welch, L. (1985). Sequential recruitment in the discrimination of velocity. *Journal of the Optical Society of America A*, *2*, 243–251.

McKee, S. P., Silverman, G. H. & Nakayama, K. (1986). Precise velocity discrimination despite random variations in temporal frequency and contrast. *Vision Research*, *26*, 609–619.

Orban, G. A., deWolf, J. & Maes, H. (1984). Factors influencing velocity encoding in the human visual system. *Vision Research*, *24*, 33–39.

van Santen, J. P. H. & Sperling, G. (1984). A temporal covariance model of motion perception. *Journal of the Optical Society of America A*, *1*, 451–473.

Simoncelli, E. & Heeger, D. (1992). A computational model for perception of two-dimensional pattern velocities. *Investigative Ophthalmology and Visual Science (Supp.)*, *33*, 954.

Smith, J. A. & Grzywacz, N. M. (1993). Model for transparent motions based on spatio-temporal filtering. In Eeckman, F. H. & Bower, J. M. (Eds), *Computation and neural systems 1992*. Dordrecht: Kluwer.

Snowden, R. J. & Braddick, O. J. (1991). The temporal integration and resolution of velocity signals. *Vision Research*, *31*, 907–914.

Tolhurst, D. J. & Movshon, J. A. (1975). Spatial and temporal contrast sensitivity of striate cortical neurons. *Nature*, *257*, 674–675.

Watamaniuk, S. N. J. & Duchon, A. (1992). The human visual system averages speed information. *Vision Research*, *32*, 931–941.

Watamaniuk, S. N. J. & Sekuler, R. (1992). Temporal and spatial integration in dynamic random-dot stimuli. *Vision Research*, *32*, 2341–2347.

Watson, A. B. & Ahumada, A. J. (1985). Model of human visual motion sensing. *Journal of the Optical Society of America A*, *2*, 322–341.

Werkhoven, P., Snippe, H. P. & Toet, A. (1992). Visual processing of optic acceleration. *Vision Research*, *32*, 2313–2329.

Yuille, A. L. & Grzywacz, N. M. (1988). A computational theory for the perception of coherent motion. *Nature*, *333*, 71–74.

Acknowledgements—This work was supported by the Smith-Kettlewell Eye Research Institute and AFOSR F49620-92-J0156. We thank Norberto Grzywacz and Suzanne McKee for their many helpful comments during the planning and execution of this project.