Behavior Research Methods & Instrumentation
1976, Vol. 8 (2), 144-151

Movement perception in computer-driven visual displays

GEORGE SPERLING
Department of Psychology, New York University, New York, New York 10003
and
Bell Laboratories, Murray Hill, New Jersey 07974

Computer-driven visual displays (CDVDs), like television and movies, produce stroboscopic rather than continuous physical movement. The success with which the perception of motion is produced depends on factors such as the fineness of the raster and the temporal and spatial relationships of the stimulus points. For a given velocity, the more points there are on the movement trajectory, and the closer their spacing, the better is the perceived movement. Moderately slow retinal velocities (on the order of .4 to .8 deg/sec) produce the highest quality of perceived movement. One can discriminate among possible subclasses of movement detectors by presenting a complex sequence of intensities at two or more points and varying their cross correlation. Motion between two areas can be perceived even when there is zero correlation between the spatial patterns in each location. Perceived motion can be of rotation, as well as of translation. The two-dimensional shadow of a rotating three-dimensional wire figure is perceived as a rotating, rigid, three-dimensional wire figure (the kinetic depth effect). A three-dimensional "shadow" of a hypothetical four-dimensional wire figure also has been produced; it was not seen as rigid.

Computers can control movements in countless ways. I shall deal here with only one: a computer generates a visual display on a cathode ray oscilloscope (CRO), and the display is viewed directly by an observer. I call this a computer-driven visual display (CDVD).

The special talents of CDVDs fall into three classes: Computers can generate complex pictures, complex trajectories, and complex sequences. I will not discuss interactive uses in which, for example, CDVDs are made contingent on eye movements.

It is often possible to produce displays similar to CDVDs in ways that do not involve computers; for example, the rotating disk with slit aperture (Figure 1). But a CDVD will generally be more versatile than the alternatives, though it may not be easy to program and it may not even be better. Digital computers with CRO displays are remarkably untalented when it comes to displaying continuous movements. Let us examine the problem.

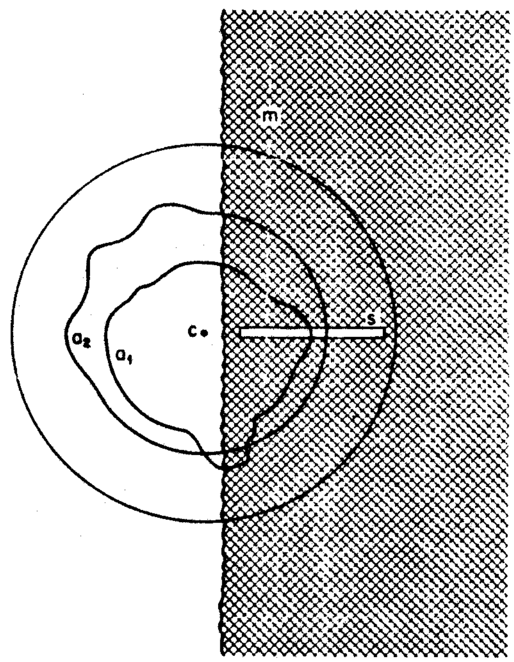

Figure 1. The aperture and rotating disk—a machine for producing translation trajectories. The area m represents an opaque mask with a narrow slit s through which the observer views the disk. As the disk rotates around its center c, the lines a_1 and a_2 appear as short moving lines in s, following whatever complex trajectories are painted on the disk.

TRANSLATION TRAJECTORY OF AN OBJECT

An object is defined as a point or group of points that remain in a fixed relation to each other. A display is composed of one or more objects. Let one point of the object be designated as the reference point. The spatial coordinates (x,y) of this point determine the translation trajectory of the object. The translation trajectory is a many-one mapping M from the time (t) into space (x,y), M:(t)→(x,y), where x and y

Presented at the National Conference on the Use of On-Line Computers in Psychology, Boulder, Colorado, November 5, 1975. The author wishes to acknowledge the helpful suggestions of Judith Harris and Saul Sternberg.

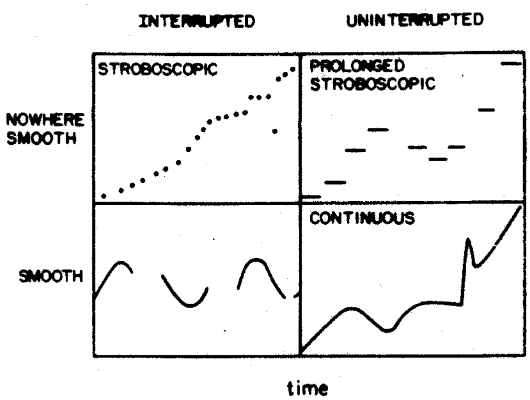

Figure 2. Movement trajectories, y(t). The abscissas represent time; the ordinates represent position, y. See text for details.

represent two-dimensional coordinates of the reference point. An equivalent way of describing a translation trajectory is by two parametric functions: x(t) and y(t). Often we say simply *a movement* to a translation trajectory. We use the word *path* to refer to the set of {x,y} values.

We say a movement is an uninterrupted movement during the time interval (a,b) if the set of t values ($a \leqslant t \leqslant b$) on which the mapping is defined has no missing segments. This means that the display is not turned off at any time during the time interval.

A movement is a smooth movement if the mapping of x(t) and y(t) is continuous in every interval on which M is defined and if the path is everywhere a connected curve, that is, it has no missing segments. A movement that is not smooth anywhere is called stroboscopic.

A smooth, uninterrupted movement is called continuous. Examples of translation trajectories are shown in Figure 2. Here, for simplicity the path is restricted to lie on a vertical line, that is, the movement is described simply by y(t).

Cathode ray oscilloscopes typically are capable of producing a continuous translation trajectory of a single point (the beam). But the way CROs normally are interfaced to digital computers makes it impossible to produce either uninterrupted or smooth movement, even of an object consisting of a single point. The significant limit here is not so much that the CDVD does not produce smooth movement (it does not), but that it produces only one point at a time. Therefore, to display an object of more than one point, production of the first point must be interrupted to produce the second, and so on.

As a practical matter, one may regard movement produced by CDVDs as stroboscopic. Points or straight-line segments are produced in a very short period of time (tens of microseconds) and repeated at a new location after intervals on the order of milliseconds or tens of milliseconds. In this respect, CDVDs are stroboscopic like cine movies and television displays. However, CDVDs can produce many more pictures per second than movies or TV, so their range of possible translation trajectories is much greater. Because light is quantal in nature, even real movement of a real object is not perfectly continuous. By producing enough discrete pictures per second, a CDVD with a stroboscopic movement trajectory can produce an arbitrarily good approximation to a continuous translation trajectory. The usual limiting factor in a CDVD reproduction of a translation trajectory is the number of spatially discrete points the CDVD can produce, that is, the fineness of the raster. With a given number of raster points available, as the approximation to a continuous translation trajectory becomes better, the maximum range of the movement shrinks. For example, suppose the fineness of the raster is 1,000 by 1,000 (in the x and y dimensions). If the resolution required is 1,000 points/deg of visual angle, the observer must stand far enough away that 1,000 points of the raster subtends only 1 deg of visual angle. Then, the CDVD's range of movement cannot exceed 1 deg of visual angle.

There are other significant problems of movie and TV displays that I will mention only briefly. In movie and TV cameras, the effective exposure is long enough to produce appreciable image blur of fast-moving objects. Blur can be avoided by using stroboscopic illumination during photography. But when a movie or TV trajectory is not a good approximation to a continuous translation trajectory, blurring the object actually makes the movement appear more realistic.

TV and CDVD displays "paint" different parts of an object at different times, and so they necessarily display different parts of a fast-moving object at different phases of the movement. The visual system is exquisitely sensitive to this kind of distortion (Ross, 1974). A related problem can occur in movie cameras that use a revolving blade or focal-plane shutter, if different parts of the object are exposed in different phases of the movement.

To reiterate, movies, TV, and CDVD display systems cannot produce a continuous translation trajectory; they can only approximate it. Figure 3 deals with uniform, linear, continuous movement and how such a translation trajectory is approximated by TV, movie, and CDVD displays. Figure 3c illustrates a TV display in which, because of persistence of the CRO phosphor, the number of objects increases to two: the first one has not faded away before the second is turned on. However, this display is perceived as only one object.

In Figure 3b, the "better" movie projector differs from the poor one (Figure 3a) in that the good projector "chops" the beam several times during a single frame to reduce the fundamental component of flicker, and it darkens the picture a smaller fraction of the time.

Figures 3d through 3h illustrate CDVD translation trajectories with rasters of various fineness and with various interintensification intervals. Figures 3d and 3g illustrate two cases where raster size Δr and interintensification interval Δt are perfectly matched. Figures 3e and 3h illustrate mismatched Δr and Δt. There is a discrete raster of points at which the object can be represented, and the real movement does not happen to pass through one of these points at precisely the scheduled time for an intensification. The usual approximation is to intensify the nearest raster point at the scheduled time, and this produces physically "jerky" movement. That is, the distance Δx between successive intensifications varies from place to place, so that velocity v ($v = \Delta x/\Delta t$) varies as Δx varies (assuming the time between intensifications Δt is held constant). For example, in 3e, the CDVD trajectory consists of three intervals with v slightly faster than the intended trajectory, and one motionless interval. This three-and-one cycle repeats. In the CDVD trajectory of Figure 3h, two motionless periods, each of three intensifications, alternate with one period of two intensifications. On the other hand, Figures 3d and 3f show perhaps equally good approximations to the same real movement. In Figure 3d, the object is intensified once at each new position; in Figure 3f it is intensified "continuously" (i.e., eight times) between translations. Whether or not these trajectories are perceived as jerky depends, of course, on many factors.

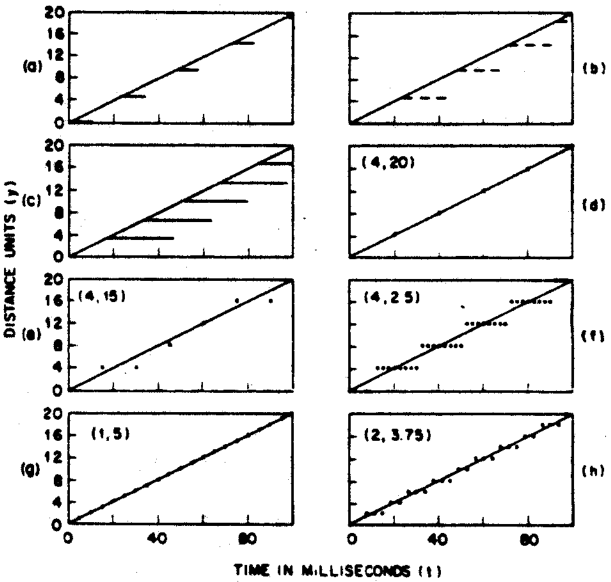

Figure 3. Constant velocity linear movement as depicted by cine movies, television, and computer-driven visual displays (CDVDs). Abscissa represents a total time of 100 msec; ordinate represents linear distance of 20 units (arbitrary scale). (a) Movie projection, once-chopped beam, 24 frames/sec; (b) Movie projection, multichopped beam, 24 frames/sec; (c) American TV, 60 half-frames/sec. The horizontal elongation of the lines represents phosphor decay time of the TV display tube; (d-h) CDVDs with various interintensification intervals (Δt) and rasters (Δr). The Δr values in units and the Δt values in milliseconds are given in the graphs as ($\Delta r, \Delta t$).

THE PERCEPTION OF MOVEMENT

There are two basically different kinds of questions that are asked about movement perception. The first is empirical: how do we perceive translation trajectories such as those illustrated in Figure 3; e.g., do the objects appear to move? to flicker? to be blurred? The second question is theoretical: what are the mechanisms of movement perception? Perhaps if we had the answer to the second question, we could answer the first, but we are far from that happy state. Attempts to answer the empirical question lead to experiments that are quite different from studies attempting to answer the theoretical question. I will deal first with the empirical question.

There is very little data on the perception of stroboscopic trajectories such as those in Figure 3 because psychologists have concentrated their researches on "simpler" displays. The most common trajectory studied is stroboscopic movement between just two points (Figure 4a), the kind of stimulus popularized by Exner, Wertheimer, Korte, Neuhaus, and hundreds of others (Boring, 1942; Hochberg, 1973). Miriam Kaplan, a graduate student at NYU, and I attempted to remedy this situation. We investigated the perception of numerous kinds of stroboscopic movement trajectories of an object composed of just one spot. Some of the trajectories we studied are illustrated in Figure 4c and 4e. We varied the space Δx between successive points on the path and the time Δt between successive exposures of the object. For each Δx and Δt, we investigated the perception of paths composed of many points and paths composed of two points.

The relationship between perceived movement, eye movement, and retinal image movement is a complex one. Movement can be perceived even when there is no change in the location of the image on the retina, as when the eye accurately tracks an object with a highly predictable trajectory. But the perceived trajectory of an object, as well as its apparent shape, is affected both by the object's motion and the eye's (cf. Hempstead, 1966). In our study, we wanted to minimize eye movements. The subjects fixated a reference point during the test and had no prior knowledge of the speed or direction of the trajectory.

Subjects made several different kinds of judgments but perhaps the most informative was that of the "quality" of the perceived movement. Subjects rated the quality of the movement on a scale from 0 to 10. A

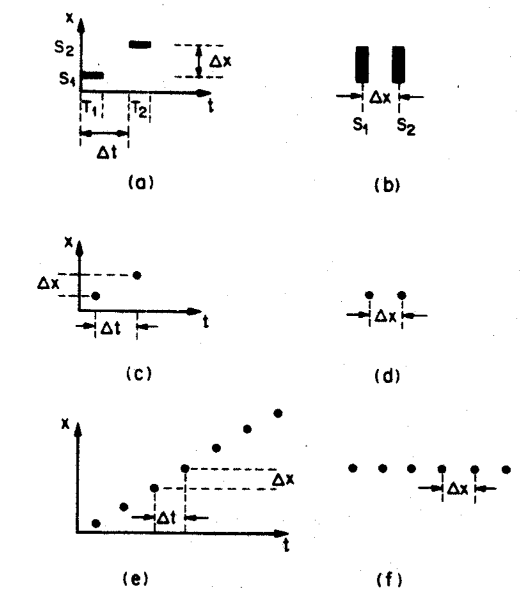

Figure 4. (a) The classical two-position stroboscopic movement trajectory. Abscissa (t) represents time, the time of occurrence of stimulus S_1 is represented by the interval T_1, and similarly for S_2. The ordinate (x) represents the spatial location of S_1 and S_2. The classical parameters varied were the distance moved, Δx, and the time from onset of the first to onset of the second, Δt. Also studied: duration of the flashes T_1 and T_2, luminance of S_1 and S_2, background luminance, configuration and size, retinal location, etc. (b) Spatial arrangement of typical, classical stimuli drawn to scale. (c) Sperling and Kaplan's trajectory for two-point stimuli and (d) their spatial arrangement; (e,f) same, multipoint stimuli.

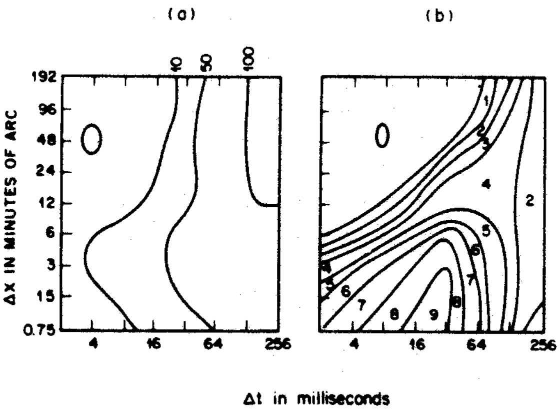

Figure 5. Judged quality of perceived movement as a function of the distance Δx and the time Δt between exposures of the points. Data of one subject; 10 judgments were made for each Δx and Δt combination. (a) Two-point experiment. The area designated zero indicates that generally no movement was perceived, i.e., less than 10% of quality judgments were greater than zero. The contours indicate boundaries of areas in which 10%-50%, 50%-100%, and lastly, in which all quality judgments were greater than zero. (b) Multipoint data. The numbers in the areas designate the median value of the quality judgments for (Δx,Δt) within the area. The quality range is 0 (no perceived movement) to 10 (apparently continuous real movement).

high rating indicated a perception similar to that of a continuous translation trajectory; a low rating indicated such defects as "jerkiness," "flicker," "object appears to turn off," etc. Zero rating meant "no perceived movement."

Figure 5a shows results with the two-point trajectories for one typical observer. The quality ratings for two-point trajectories are so low that it is more informative to indicate, not the average rating, but simply the proportion of trials that yielded a nonzero rating, i.e., a minimal perception of movement. Our data show that large Δx and Δt yield the most reliable perceptions of motion. In fact, our data are quite similar to those of Neuhaus (1930). We know that if we had increased Δt or Δx beyond the maximum values illustrated, quality would have diminished.

Results with many-point displays are shown in Figure 5b. The most obvious result is that the quality ratings are enormously higher, and the range of Δx and Δt giving rise to good motion perception is enormously greater. It is not apparent from Figures 5a and b, but the quality of perceived movement in every two-point display was improved by adding to it more points with the same Δx and Δt. A more important result is that for a given velocity, $v = \Delta x/\Delta t$, the greater the density of points along the path (i.e., the smaller Δx and Δt), the higher the quality of the perceived movement. With a sufficient density of points on the path, it was possible to produce stroboscopic movement of nearly the same quality as continuous movement.

The two-point display data might have led us to design displays with large Δx and Δt in order to produce good apparent movement. The practical conclusion is that the smaller are Δx and Δt, the higher is the quality of perceived movement.

These data apply to a single-point object and a particular class of stroboscopic trajectories, i.e., uniform linear translation. We have not yet generalized the point results to larger objects (composed of many points), nor even predicted the multipoint data from two-point data. Perceived movement is a complicated phenomenon, and when we look at theoretically oriented experiments we shall see why.

MODEL OF A TRANSLATION TRAJECTORY DETECTOR

A simple model of a translation detector (Figure 6) contains three kinds of components (e.g., see Reichardt, 1961): (1) Two input transducers receive visual input from two different areas (A_1 and A_2); they transmit their outputs $f_1(t)$ and $f_2(t)$ to (2) a comparator, which compares f_1 and f_2 with each other and produces a signal g(t) that indicates their similarity; (3) A detector D receives inputs g(t) from one or more comparators and perhaps also from some other sources, such as flicker detectors; D outputs a categorical "detection" response, corresponding to the perception of motion.

One additional complication at this point will simplify matters later. Consider a detector for movement in direction A_1 to A_2 and an object moving in the direction A_1 to A_2. If we assume the comparator has no memory, then signals passing through the A_1 channel must travel more slowly than signals through the A_2 channel in order that the object-induced signals will reach the comparator and be compared simultaneously. If we designate the transduction operation of the A_2 channel as T, then we can designate the transduction operation of the A_1 channel as $T + \tau$, where τ is a "delay" operation with a mean delay time $\bar{\tau}$.

One can ask many questions about such a model. What is the spatial separation between the areas? What is their spatial extent? What is the nature of the transduction and of the comparison operation? What is the mean time delay $\bar{\tau}$ of f_1 with respect f_2? For example, the ratio, $\Delta x/\bar{\tau}$, gives the velocity of movements to which this detector is very sensitive. The sizes of Δx, A_1, and A_2 determine the size of the detector, i.e., the area within which it detects movement. The nature of T determines the kind of stimuli that elicit a movement response.

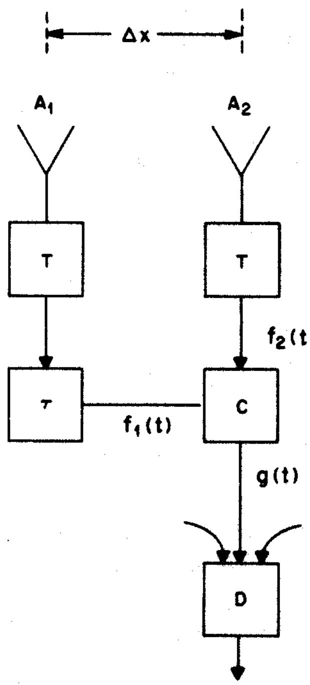

Figure 6. Model of a movement detector. A_1 and A_2 are input areas on the retina, T is a transformation of the received input, τ is an additional "delaying" operation, C is a comparator, and D is the detector.

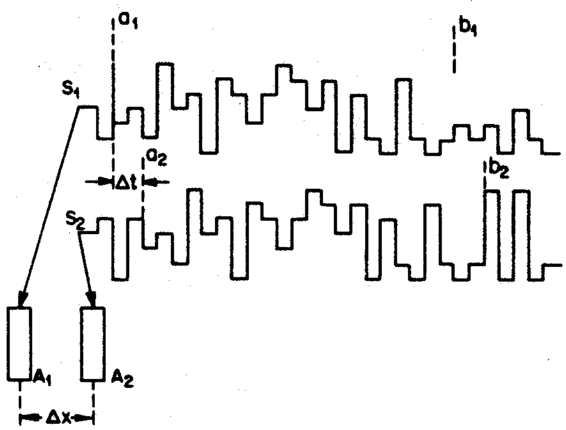

Figure 7. Method of complex sequences for isolating subclasses of motion detectors. S_1 and S_2 represent complex sequences of luminance (ordinate) as a function of time (abscissa) presented in rectangles A_1 and A_2, respectively. In the example, sequence S_2 repeats a part (a_1,b_1) of sequence S_1 after a time lag, Δt.

Of course, the movement detector of Figure 6 is not the only conceivable kind. One interesting question is whether or not human movement detectors actually resemble this model. A system based on detectors of this kind would, for example, require at least two detectors with different parameters to discriminate between different velocities of movement. Another kind of detector might discriminate different velocities all by itself. For humans we know few answers. We have better information about insect motion perception, and for an account of progress in this domain I again refer the reader to the pioneering work of Reichardt and his collaborators (Reichardt, 1961).

Given the two-input single-delay model of Figure 6, it is easy to see why one would be led to do the classical two-point stroboscopic movement experiments (Figure 4a). This would be an ideal experiment if there were just one kind of movement detector—just one set of parameters. When there are many detectors, each with a different Δx and τ, this experiment does not discriminate among them: it only gives the properties of the whole aggregate. In fact, the two-point experiment is one of the worst for discriminating among detectors.

To obtain more detailed information about subclasses of motion detectors, a more complicated experiment is needed; it is here that the CDVD becomes useful. Miriam Kaplan and I have found one kind of experiment that can provide detailed information about subclasses of detectors (Note 1). It consists of presenting complex sequences of stimuli at each of two locations, and studying the perception of motion as a function of the properties of the sequences, particularly their cross correlation (Figure 7). The stimuli are bars of light which fluctuate in luminance. When the pattern of luminance changes in one bar is

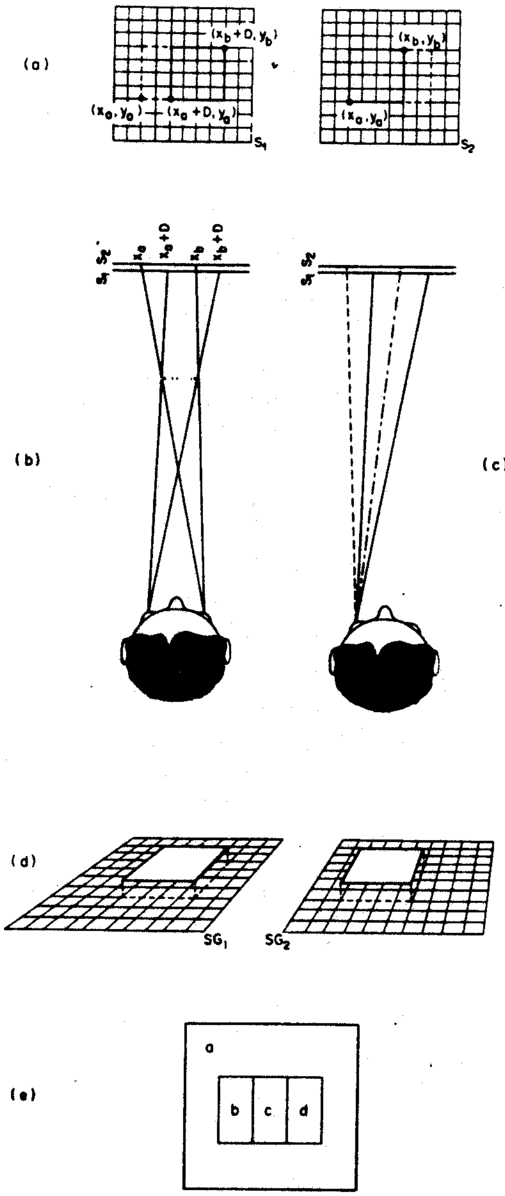

Figure 8. Random dot stimuli for depth and motion perception. (a) Construction of a stereogram. The critical area of stimulus S_1 (square with lower left corner $[x_a + D, y_a]$) has a texture identical to the corresponding area of S_2 (square with corner $[x_a, y_a]$). (b) How the depth illusion is generated in a stereogram. S_1 is seen only by the left eye, S_2 by the right; they are at the same physical depth but shown separately for ease of depiction. The dotted line indicates the apparent extent of the critical area; its distance from the background indicates the apparent depth of the critical area relative to the background. (c) Consecutive exposure of S_1 and S_2 yields apparent motion of the critical area. Monocular viewing is illustrated; binocular viewing yields the same perception. (d) SG_1 illustrates the depth perception produced by a complete stereogram; similarly for SG_2. The critical area is translated laterally by distance D in S_2 relative to S_1. When viewed consecutively as in (c), the critical area appears to move laterally. (e) Movement without correlation in a random dot stimulus. Area a remains unchanged, b changes on even frames, d on odd frames, c on all frames. Area b + c appears to move to c + d and vice versa.

reproduced after a certain brief interval in the second bar, an illusion of movement results. The illusion may be likened to a train of boxcars of various brightness passing behind and being seen through two windows.

The advantage of this method of complex sequences is that it is possible to produce sequences in which: (1) correlation exists only for a particular temporal (Δt) and spatial (Δx) separation, and (2) there is a large amount of "noise" at other separations. The noise prevents a detector from responding to a signal with a particular Δt and Δx, unless the detector is exactly tuned to Δt and Δx. Thus the method permits the psychophysical isolation and measurement of movement detectors. This research is still at an embryonic stage, but it is a good example of a complicated experiment that works easily and well on CDVDs.

COMPLEX MOVEMENTS AND OBJECTS

It is in the domain of complexity that the CDVD begins to live up to its potential. Two kinds of complex displays will be considered. The first is the "random dot stereogram" (Julesz, 1971). The stereogram consists of two monochromatic stimuli, S_1 and S_2, which, when combined stereoscopically, give rise to the perception of depth—one portion appears to be in front of or behind the rest. S_1 and S_2 can be described by their luminance distributions, $\mathcal{L}_1(x,y)$ and $\mathcal{L}_2(x,y)$. Each stimulus is divided into little areas $(\Delta x, \Delta y)$; typically there is a 100×100 grid, producing 10,000 squares. Each little area $(\Delta x, \Delta y)$ of S_1 is painted a luminance $\mathcal{L}_1(x,y)$, where each value of $\mathcal{L}_1(x,y)$ is chosen randomly from a set of possible luminances. The luminance distribution $\mathcal{L}_2(x,y)$ is taken to be equal to $\mathcal{L}_1(x,y)$ except in a "critical" area and two adjacent areas. In the example (Figure 8), the critical area appears to be out in front. The adjacent areas correspond to those parts of the background that would be hidden (occluded) from the view of one eye or the other by the critical area, if it were really out in front. Suppose the critical area is a rectangle defined between x_a and x_b and between y_a and y_b; that is, $x_a \leq x \leq x_b$, $y_a \leq y \leq y_b$ (Figure 8a). In this area, the luminance $\mathcal{L}_2(x,y)$ is chosen such that $\mathcal{L}_2(x,y) = \mathcal{L}_1(x + D,y)$. That is, the pattern of light and dark squares within the rectangle is reproduced an amount D to the left in S_2, relative to S_1. Luminance values in the occluded rectangle of S_2 (between x_b,y_a) and $(x_b + D,y_b)$ are chosen randomly, independent of S_1. The net result can be interpreted in terms of operations on two identical pieces of textured wallpaper, S_1 and S_2. A rectangular piece of S_2 is cut and moved left; the space remaining is filled with a new, unrelated piece of texture.

When the stereogram of Figure 8a is viewed by an observer with S_1 seen by the left eye and S_2 by the right, the rectangle is perceived as standing out in front of the background (Figure 8b). When S_1 and S_2 are viewed successively, i.e., S_1 is presented first to either or both eyes and replaced by S_2 a fraction of a second later in the same location, then the rectangle is perceived as moving leftward by an amount D (Figure 8c). In fact, the rectangle does have a translation trajectory typical of those used to study apparent movement (Figure 4a). What makes this stimulus fundamentally different from that of 4a, however, is that in 4a the critical area (rectangle) can be discriminated from the background (blank) in either stimulus alone; in the Julesz stereogram, the critical area of stimulus S_1 is defined only in relationship to stimulus S_2—until both S_1 and S_2 are known, the critical area cannot be defined.

There are many experimental variations of the binocular paradigm in which the parameters of S_1 and S_2 are varied. For each of these binocular variants, there is a corresponding motion variant in which the two stereograms are exposed successively rather than simultaneously (cf. Figures 8b and 8c), and the perception of motion occurs instead of the perception of stereoscopic depth (Anstis, 1970; Julesz, 1971; Ramachandran, Madhysudhan Rgo, & Vidyasagar, 1973; White, 1962).

Even more complex designs can be created by combining the depth and motion paradigms. For example, Julesz and Payne (1968) presented series of consecutive stereograms in which the critical area was produced at a certain depth in each stereogram but translated laterally between successive stereograms (Figure 8d). The observer correctly perceived the critical area as moving.[1] This is a theoretically interesting result. In terms of the kind of model shown in Figure 6, it means that the input to areas A_1 and A_2 can come not only from the retina but also from a binocular projection field in which the inputs from the two retinas are already combined (cf. Sperling, 1970).

One perplexing observation I made recently complicates the interpretation of the above experiments: To see motion it is not necessary that the critical areas of S_1 and S_2 be identical. Motion phenomena can be observed even when, in the critical area, $\ell_1(x,y)$ and $\ell_2(x,y)$ are absolutely uncorrelated, i.e., are different and have no systematic relation to each other. This paradigm is illustrated in Figure 8e. The stimulus S is composed of a random texture that is divided conceptually into four areas as illustrated. Area a remains unchanged. The texture in other areas is repeatedly replaced with new random textures according to the following schedule: c changes on every frame, b changes on odd frames, and d changes on even frames. When stimulus frames are changed every 50-150 msec, an observer perceives the rectangle defined by areas b + c as moving to and becoming the rectangle defined by areas c + d on even frames and then moving back on odd frames. Other perceptions of movement are also possible, such as area b moving to become area d, and vice versa. As the duration between frames lengthens, this experiment can become another measure of visual persistence ("iconic memory"), cf. Sperling (1960, 1963).

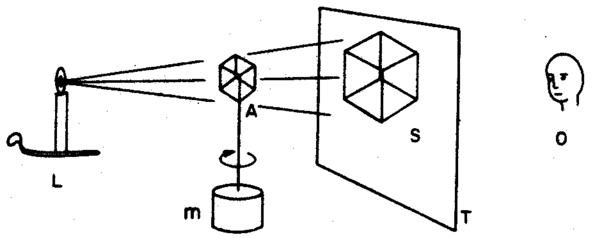

Figure 9. Demonstration of the kinetic depth effect. A point source of light L casts a shadow S of an object A. The object is rotated by a motor M and its shadow is viewed from behind a translucent screen T by the observer O.

KINETIC DEPTH EFFECT

Perhaps the most interesting movements are not translations at all, but rotations.

An object consisting of more than one point may undergo rotations and translations. Here the usual perceptual question does not concern the quality of perceived movement directly, but rather whether the object is seen as rigid when it is undergoing rotation. We certainly perceive the ordinary, solid, three-dimensional objects of everyday experience as remaining rigid when they move. There is a redundancy of cues.

What happens when we eliminate some cues to solidity? An easy way to do this is to view the shadow of a rotating three-dimensional object rather than the object itself. In many cases, the shadow contains sufficient cues so that it is correctly perceived as the shadow of a rigid three-dimensional object rather than as a nonrigid two-dimensional "object," i.e., a two-dimensional object that changes its shape as it moves. The phenomenon, first studied systematically by Wallach and by Gibson and their respective collaborators (Gibson & Gibson, 1957; Wallach & O'Connell, 1953), is called the *kinetic depth effect: kinetic* because the shadow must move, and *depth* because when the shadow does move, the flat two-dimensional shadow is perceived as representing a three-dimensional object having depth.

The kinetic depth effect works especially well with wire figures such as a three-dimensional wire cube (Figure 9). This observation that a three-dimensional wire cube can be perceptually reconstructed from its two-dimensional shadow intrigued my colleague, A. M. Noll, at the Bell Labs. He asked, "Can a four-dimensional hypercube be perceptually reconstructed from its three-dimensional shadow?" This problem is perhaps the ideal one to demonstrate the power of the CDVD.

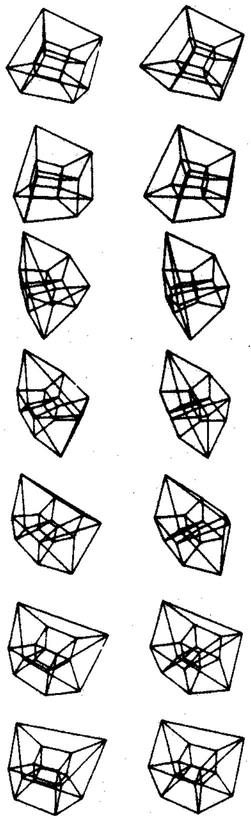

Figure 10. Two-dimensional shadows of a four-dimensional wire hypercube. The left and right columns represent the left eye's and the right eye's views of a dynamic stereogram. Successive rows represent successive instants in time during a rotation of the hypercube (from Noll, 1967a).

The extension of projective transformations to four dimensions and the resulting computations of rotation are not practical without computers. The three-dimensional "shadow" is achieved by presenting two two-dimensional projections, one to each eye, to produce a three-dimensional stereoscopic depth effect (Noll, 1967b; cf. Sperling, 1971, for stereoscopic methods; White, 1962).

Figure 10 shows some two-dimensional shadows (the left and the right eye's views) of the four-dimensional hypercube. Each pair of pictures, individually, is perceived as a three-dimensional wire object. Has anyone viewing these changing projections ever had a perception of a solid four-dimensional object? Not yet—certainly not this observer. Is this failure due to a lack of practice, or is it due to some fundamental structural limit of the human mind? We do not yet know. But it is to answer fascinating questions like this that the CDVD becomes not merely a convenience but an absolute necessity.

REFERENCE NOTE

1. Sperling, G., & Kaplan, M. Unpublished experiments conducted at New York University, 1974-6.

REFERENCES

Anstis, S. M. Phi movement as a subtraction process. *Vision Research*, 1970, **10**, 1411-1430.

Boring, E. G. *Sensation and perception in the history of experimental psychology*. New York: Appleton-Century-Crofts, 1942.

Gibson, J. J., & Gibson, E. J. Continuous perspective transformations and the perception of rigid motion. *Journal of Experimental Psychology*, 1957, **54**, 129-138.

Hempstead, Charles F. Motion perception using oscilloscope display. *IEEE Spectrum*, 1966, **3**, 128-135.

Hochberg, J. Perception. II. Space and movement. In J. W. Kling & L. A. Riggs (Eds.), *Experimental psychology* (3rd ed.). New York: Holt, Reinhart, & Winston, 1973. Pp. 475-550.

Julesz, B. *Foundations of cyclopean perception*. Chicago and London: The University of Chicago Press, 1971.

Julesz, B., & Payne, R. A. Differences between monocular and binocular stroboscopic movement perception. *Vision Research*, 1968, **8**, 433-444.

Neuhaus, W. Experimentelle Untersuchungen der Scheinbewegung. *Archiv für gesamte Psychologie*, 1930, **75**, 315-458.

Noll, A. M. Computers and the visual arts. In *Design and planning, Number 2*. New York: Hastings House Publications, 1967. Pp. 65-79. (a)

Noll, A. M. A computer technique for displaying n-dimensional hyperobjects. *Communications of the ACM*, 1967, **10**, 469-473. (b)

Ramachandran, F. S., Madhysudhan Rgo, V., & Vidyasagar, T. R. Apparent movement with subjective contours. *Vision Research*, 1973, **13**, 1399-1401.

Reichardt, W. Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In W. A. Rosenblith (Ed.), *Sensory communication*. Cambridge, Mass: M.I.T., and New York, London: Wiley, 1961.

Ross, J. Stereopsis by binocular delay. *Nature*, 1974, **248**, 363-364.

Sperling, G. The information available in brief visual presentations. *Psychological Monographs*, 1960, **74**, No. 11 (Whole No. 498).

Sperling, G. A model for visual memory tasks. *Human Factors*, 1963, **5**, 19-31.

Sperling, G. Binocular vision: A physical and a neural theory. *American Journal of Psychology*, 1970, **83**, 461-534.

Sperling, G. Stereoscopic visual displays: Principles, viewing devices, alignment procedures. *Behavior Research Methods & Instrumentation*, 1971, **3**, 154-158.

Wallach, H., & O'Connell, D. N. The kinetic depth effect. *Journal of Experimental Psychology*, 1953, **45**, 205-217.

White, B. W. The computer as a pattern generator for perceptual research. *Behavioral Science*, 1961, **6**, 253-259.

White, B. W. Stimulus-conditions affecting a recently discovered stereoscopic effect. *American Journal of Psychology*, 1962, **75**, 411-420.

NOTE

1. Julesz and Payne also observed a phenomenon they called "binocular standstill." This probably is not a binocular phenomenon at all—merely a variant of the effect illustrated in Figure 8e when viewed with both eyes.