

How to Study the Kinetic Depth Effect Experimentally

George Sperling
New York University

Barbara A. Doshier
Columbia University

Michael S. Landy
New York University

Sperling, Landy, Doshier, and Perkins (1989) proposed an objective 3D shape identification task with 2D artifactual cues removed and with full feedback (FB) to the subjects to measure KDE and to circumvent algorithmically equivalent KDE-alternative computations and artifactual non-KDE processing. (1) The 2D velocity flow-field was necessary and sufficient for true KDE. (2) Only the first-order (Fourier-based) perceptual motion system could solve our task because the second-order (rectifying) system could not simultaneously process more than two locations. (3) To ensure first-order motion processing, KDE tasks must require simultaneous processing at more than two locations. (4) Practice with FB is essential to measure ultimate capacity (aptitude) and, thereby, to enable comparisons with ideal observers. Experiments without FB measure ecological achievement—the ability of subjects to extrapolate their past experience to the current stimuli.

Our article (Sperling, Landy, Doshier, & Perkins, 1989, henceforth, the *source article*) proposed the following: (1) An objective task that involves 53 different shapes to measure shape identification performance in kinetic depth effect (KDE) experiments; (2) an algorithm for the structure-from-motion computations that subjects perform on these and similar stimuli; and (3) a distinction between three kinds of computations. We distinguished (a) the true KDE computation, (b) a KDE-alternative computation—an informational equivalent to the true KDE computation but carried out elsewhere in the brain, and (c) an artifactual computation that arrives at the correct response in a given task but is based on an incidental property of the display. We motivated our discussion by pointing out difficulties in previous work on KDE that we believed could be remedied by measuring objective performance in tasks like ours.

Braunstein and Todd, two experimenters who felt unjustly criticized, wrote a commentary (Braunstein & Todd, 1990) in which they argued that (a) we dismiss legitimate 2D relative velocity cues to KDE as artifactual, (b) our experimental task was not exempt from the criticisms we levied at others, and (c) KDE should be measured in tasks in which subjects are not given feedback about the correctness of their responses (so that the subjects do not learn to use artifactual cues).

Braunstein and Todd's (1990, henceforth, *the critics*) point (a) reflects an unfortunate misreading. In fact, we proposed that "the structure-from-motion algorithm. . . involves finding local 2D velocity minima and maxima and assigning depth values to these locations in consistent proportion to their

velocities" (Sperling et al., 1989, p. 839; see also Figure 6, p. 838). In this article, we attempt to further clarify the role of relative motion cues in KDE tasks (using 2D dynamic images to answer questions about 3D shape) and in the continuum of mental computations between true KDE and truly artifactual computations. KDE is the perceptual experience of 3D object depth evoked by dynamic 2D images. The opposite end of the continuum is an artifactual computation that arrives at the correct response for the experimental task by using an incidental property of the display. Although artifactual computations need not involve motion cues, in the source article we gave examples of some that do. These computations were called artifactual because motion entered the computation in a way that shortcut the KDE computation. In one example, a measurement of absolute velocity at a single fixed point in the display would have sufficed to yield the correct response. Our aim was not so much to classify computations but to actually deduce the minimal computation that would suffice to solve particular KDE tasks. In a well-constructed task, there is no computational shortcut—the minimal computation is the KDE computation or is essentially equivalent to it.

The critics' arguments (b) and (c) were anticipated and considered in the source article. Here we elaborate the source article's discussion and respond to two newly raised fundamental issues (How should experiments be conducted? and How can a subject's mental computations be exposed, measured, and controlled?) and to other issues (immediacy, practice, and scintillation) that pertain to our specific task.

The 53-Shape KDE Task

Motion Produces an Immediate Experience of KDE

Our task uses 53 different shapes whose surfaces consist of random-dot textures. Each shape is defined by three equally spaced locations. Each location contains either a hill (+1) or

Preparation of this article was supported by the United States Air Force Life Sciences Directorate, Visual Information Processing Program Grant 88-0140.

Correspondence concerning this article should be addressed to George Sperling, Psychology Department and Center for Neural Sciences, New York University, 6 Washington Place, Room 980, New York, New York 10003.

a valley (-1), or is flat (0). Stimuli are constructed on one of two such sets of locations. A smoothing (spline) transformation merges two hills into a ridge, and merges a three-hill configuration into one spread-out mound. The static random-dot surface has no shape cues and looks utterly flat. When the surface begins to move in a gentle rotary oscillation, it appears instantly to have a particular shape. In the source article, we wrote "All subjects reported that they perceived a 3D surface the first and every time they viewed the high numerosity displays" (p. 830). We have demonstrated these KDE shapes to subjects, to visitors, and to hundreds of observers at numerous lectures and have not yet received even one report of an observer who did not perceive a vivid 3-D shape. We believe this is a natural, ecologically valid test of the shape identification functions that KDE has evolved to perform.

Effects of Practice

Under normal viewing, no learning or practice is needed to perceive these 3-D shapes. The critics' notion that practice is needed to perform the basic task is wrong. However, practice is helpful for some aspects of the task. (a) Subjects must learn to correctly use the naming convention for these shapes. This is usually learned in just a few views of sample stimuli. (b) All subjects remember the shape of the oscillating object, but unpracticed subjects frequently forget the final direction of rotation. With corrective feedback, they learn to report both shape and motion. (c) Some shapes are deliberately made quite similar. For example, in these stimuli, two adjacent hills combine to form an elongated hill. The distinction between a two- and three-hill configuration might be overlooked by a subject who did not receive feedback of the correct response, but the distinction is easily learned. However, even highly practiced subjects, with feedback, cannot infallibly discriminate between the two differently oriented three-hill configurations. (d) Small amounts of image degradation are easily tolerated by all subjects. However, to correctly identify shapes when the number of surface dots is grossly reduced or when the signal-to-noise ratio is reduced takes practice.

Feedback or No Feedback?

The critics suggest that KDE experiments be conducted without informing the subject about the correctness or incorrectness of the response (the 3D structure derived from motion), that is, giving the subject no feedback. By eliminating feedback, the critics hope to avoid the problem of the subjects' learning to use incidental cues. We believe that the better way of dealing with incidental cues is to eliminate them, or when that is not feasible, to mask them or render them useless by irrelevant variation (see below). What we address here is the larger question of what can be learned from experiments with and without feedback.

An experiment without feedback is essentially an epidemiological investigation. It investigates the current status of a skill that has been acquired prior to the experimental situation. Therefore, the no-feedback experiment is simplest to interpret when the current test is identical to a previous

learning situation. As the experimental test stimuli diverge from the original training stimuli, the experiment must be interpreted in terms of the ease of generalization of previously learned skills to the new testing stimuli. Thus, to test the ability of pilots to discern subtle terrain features in brief glimpses, we would test them with images of natural terrain. Testing pilots with our 53 random-dot shapes, without first affording them an opportunity to practice with feedback, would tell us only how their previously acquired skill generalized; it would not be appropriate for measuring either their previously learned KDE skill or their ability to acquire new skills. Feedback experiments teach us what humans can and cannot learn to do—the limits of human performance. Because experiments with feedback probe the limits of performance, they are the most informative for the discovery of processing mechanisms. In no-feedback experiments, the unknown training situation, and the divergence of test and training, pose problems for theoretical analysis.

Ideal Observers

One kind of investigation that has been particularly informative about human computation is the comparison of human performance with the performance of a statistically ideal observer (Green & Swets, 1966/1974; Sperling & Doshier, 1986). Indeed, the efficiency of information use by humans relative to ideal observers is of practical as well as theoretical interest. Tracing information loss through the stages of sensory analysis yields important insights into sensory processing (e.g., Geisler, 1989; Parish & Sperling, 1987). It would be of great interest to know, in noise-perturbed KDE displays, what the efficiency of human shape identification is in relation to that of an ideal observer. When the efficiency of human perceptual processing is high, we suspect that the task exposes processes that are of evolutionary significance.

To compare the processing efficiency of human and ideal observers, we need to specify exactly what each kind of observer knows about the experimental procedure, such as the a priori knowledge about the probabilities of various kinds of stimuli and the payoffs. This implies an experiment with explicit feedback. To test and evaluate sophisticated models of human mental computation requires us to bring into the laboratory much of the training that often is assumed to have occurred naturally, and it requires more complex and more explicit laboratory procedures than have been used in the past, all with feedback to the subject.

Introspection Versus Objective Measures of Performance

Early psychologists such as Wundt (1905) and James (1890) attempted to distinguish the new discipline of experimental psychology from the natural sciences by emphasizing different methodologies. They were especially concerned with how things appeared to them—introspection—rather than with measurable skills and abilities. An important component of the development of psychology has been the move away from introspection—now viewed as an extended verbal report—and toward behaviors that are simpler, more easily measured,

and more directly related to evolutionary development. Behavioral investigations of KDE tend to require more work than the corresponding introspective investigations, but that is a tax to which perceptionists have become accustomed. From a formal point of view, the subject's report in an introspective procedure "tell me what you see" and in an experiment without feedback "is what you see a circle or a cylinder?" have in common that they report the subjective appearance of the world unconstrained by feedback of objective reality, and they differ primarily in the degree of constraint on the response. In studying KDE, which is a critical component of a structure-from-motion system that has evolved to meet a biological demand, it is important not to stop at the point of studying appearances (perceptions) but to continue to investigate how these perceptions ultimately govern performance.

KDE, Alternative Computations, and Artifactual Computations

Motion Input to KDE Computation

A stationary surface covered with random dots appears quite flat. As soon as the surface starts to rotate, it is perceived as having depth—the kinetic depth effect (KDE). It is useful to think of KDE in relation to stereopsis, which is the perception of depth induced by the disparity differences between images in the left and right eyes. In perfect analogy to stereopsis, the disparity differences between dots in two successive frames of our random-dot displays suffice to give subjects a good impression of KDE depth (Landy, Doshier, Sperling, & Perkins, 1988; Landy, Sperling, Perkins, & Doshier, 1987; Todd, 1988). Because only velocity is computable from two frames, and not acceleration (which requires three), KDE perceived in two-frame displays implies that acceleration is not needed for KDE. Similarly, constructing a sequence of frames so that each individual dot has a lifetime of only two frames yields high shape identification accuracy (Doshier, Landy, & Sperling, 1989). These, and other of our results, imply that a velocity flow-field is a sufficient stimulus for KDE. The identities of the moving dots are preserved only long enough to yield a velocity estimate; subsequently, only the velocities and not the dots themselves are needed for the KDE structure-from-motion computation.

The human perceptual system uses two fundamentally different computations to extract motion flow-fields. First-order motion analysis is served by motion detectors that approximate a spatiotemporal Fourier analysis based on stimulus contrast (Adelson & Bergen, 1985; van Santen & Sperling, 1984, 1985; Watson & Ahumada, 1985). Second-order (nonFourier) motion analysis uses more complex stimulus properties and invariably involves rectification (e.g., the absolute value of contrast; Chubb & Sperling, 1988, 1989). Doshier et al. (1989) showed that the first-order system was the predominant contributor to KDE in our random-dot stimuli (because only it had sufficient spatial resolution), although second-order motion computations could yield limited KDE under special conditions.

Structure-From-Motion Computation

Given that the first-order velocity flow-field is the input stimulus for KDE, we specify a computation that would be sufficient to extract 3D shape from our 2D stimuli. Considering that the stimuli were viewed by parallel projection and with the axis of rotation perpendicular to the viewing axis, the principle is the same as for motion parallax: The difference between local object depth and the mean object depth is proportional to the difference between the local image velocity and the overall image velocity. The sign of the proportionality (+ or -) is undetermined. To discriminate among our 53 stimuli, it is not necessary to compute depth everywhere in the image. For example, it suffices to determine depth (velocity) at the six locations in which a 1, 0, or -1 could be placed during stimulus construction. Alternatively, it would be sufficient to locate the velocity minimum or maximum, to derive a three-valued descriptor of the extremum, and to interpolate smoothly to the reference plane everywhere else. (A plausible peak descriptor would take the values *normal* [single peak used in construction], *elongated* [the combination of two peaks], or *enlarged* [the combination of three peaks].) Although the precise location of the extremum would (in the absence of noise) be sufficient for an ideal detector to discriminate between the 53 shapes, human observers use the actual shape in the neighborhood of the extremum in their judgment.

KDE Versus Non-KDE Computations

In the source article, we proposed a continuum of computations ranging from a true KDE computation, in which 3D shape assignment is accompanied by the introspective impression of depth structure, to artifactual computations. Artifactual computations can be eliminated by appropriate stimulus manipulations. The more troublesome possibility is a *KDE-alternative computation* that is algorithmically equivalent to the true KDE computation and may share the same motion inputs, but is not accompanied by the introspective impression of a shape in depth. In the source article, we provided a task in which the subject viewed six isolated patches in which dots moved at the velocity of six key locations in our KDE displays. Performance in this task, which did not involve KDE, demonstrated not only that an alternative computation could occur, but that the pattern of responses and errors produced by the alternative computation mirrored the response pattern in the KDE task.

Here we consider perhaps the most troublesome possibility: true KDE supplemented by other computations. Suppose, for example, in viewing one of our complex random-dot shapes, the subject (a) experiences weak KDE and sees a hill and an ambiguous area, (b) observes that dots in the hill and in the ambiguous area of the display are moving in the opposite directions, (c) infers that these two subareas represent the opposite depth planes, and (d) uses both sources of information in his response. How can one deal with the problem of discovering the algorithms underlying KDE performance in KDE tasks and their precise implementation? The critics assume that subjects naturally use KDE in KDE tasks, and that by not giving subjects feedback on the correctness of their

responses, subjects will not learn to use artifactual or alternative computations.

Because subjects continuously practice natural KDE tasks in everyday life, which provides feedback, they may normally use conscious or unconscious non-KDE computations to help them to derive structure from motion. Non-KDE computations are not necessarily restricted to the laboratory, and it is worthwhile to develop methods to measure them. Because KDE is distinguished from a KDE alternative by a subjective impression of perceived depth, subjective reporting is a necessary component in the perceptionist's arsenal. By clever experimentation, artifactual computations usually can be distinguished from KDE and KDE alternatives. Whereas the critics rely primarily on subjective reports and introspection in approaching the problem of artifactual and alternative computations, we feel it is necessary to augment the introspective approach with successive refinements in experimental procedures to gain control over critical factors as they are discovered. For example, in our shape identification task, using shapes of considerable complexity makes artifactual computations (which rely on small amounts of incidental information) almost useless. Keeping viewing times fairly brief create a relative disadvantage for serial (versus parallel) computations. To further discriminate between KDE and non-KDE computations, the source article proposed dual tasks that selectively interfered with the mental processing resources required for alternative or artifactual computations.

Multilocation Motion Tasks

The source article argued, and the critics appear to agree, that the 53-shape identification task is less susceptible to artifactual computations than earlier tasks. However, the critics argue that the distinction between the 53-shape task and previous KDE tasks is not based on any fundamental principle, but is based merely on the number of locations from which velocity information must be extracted in order to perform the task. In rebuttal, we show here that the distinction between one or two versus six locations is critical (because either the first- or second-order motion system can provide simultaneous velocity information about one or two locations, but only the first-order system can provide information about six). In the source article, we showed that the structure of the earlier tasks permitted information from only one or two locations to discriminate perfectly between alternative responses when the same information would have been insufficient to construct a 3D shape representation.

Dosher et al. (1989) showed that complex shape identification, based on motion at three or more locations, operates very differently from motion extraction at one or two locations. They used stimuli that were designed to selectively stimulate either the second-order motion system or both the first- and second-order systems and compared them in four tasks: the 53-shape identification task, a threshold detection task for motion in a single patch, a threshold direction-of-motion task in a single patch, and a motion segmentation task that required finding the one-of-nine possible locations at which there was an odd direction of motion. Manipulations that disrupted first-order motion information (such as rapidly alternating the contrast of stimulus dots on a gray ground

between black and white) reduced identification of 3D shapes to near chance levels. Such manipulations also severely impaired the motion segmentation task. Performance was equivalent to sophisticated guessing that uses velocity information based on only one or two locations. In contrast, detection performance and direction-of-motion judgments in single patches of planar motion survived disruption of first-order information.

There are several reasons why artifactual computations based on one or two locations may survive first-order motion disruption when KDE computations cannot. Velocity information about one or two locations may be obtained either by tracking individual image features, or by information processed through a second-order system. However, second-order motion computations are based on some form of rectification (either half wave or full wave) that implies loss of information in relation to first-order motion (Chubb & Sperling, 1988; Sperling, 1989). Empirically, second-order information is virtually restricted to the fovea, and even there it is of low spatial resolution (Chubb & Sperling, 1988, 1989). Not surprisingly, when the recovery of 3D shape requires simultaneous information about motion in three or more locations, it is extremely vulnerable to disruption of first-order information.

The Dosher et al. (1989) results are one cogent empirical example of why we make the distinction between tasks that require observation of only one or two locations and those that require more. The 53-shape KDE task has been demonstrated to require the first-order motion system. Indeed, first-order motion appears to be the essential input to all complex KDE discriminations. Tasks that can be solved by extracting velocities at two locations do not necessarily involve either KDE or the first-order motion system. Our criticism of other tasks has been that they easily can be solved by computations that are less than the whole KDE computation. Indeed, it would be a step forward if the KDE experimenters offered a plausible KDE computation against which performance in their tasks could be measured.

How to Deal with Artifactual Cues

Dot Density: An Artifactual Cue in KDE Experiments

The source article investigated structure from motion. In this context, a nonmotion shape cue such as a local variation in 2D dot density is artifactual. (Of course, in a shape-from-texture task, texture-density cues would be primary.) We eliminated static (single-frame) density cues by resampling of stimulus dots when necessary to maintain a uniform 2D density on each successive frame of the motion stimulus, as follows. The stimulus field was divided into 100 fixed areas of equal size. Each area was constrained to have three dots in every frame. The 3D motion of the surface between frames causes dots to wander in and out of areas, some areas having a net inflow and others a net outflow. Therefore, to satisfy the constraint of having a constant number of dots, dots were added or subtracted at randomly chosen locations within local areas. The fraction of new-plus-discontinued dots divided by the fixed number of dots in an area is the scintillation fraction. Our displays typically required 5% frame-to-frame scintillation.

In the source article, we measured the density artifact by producing a display that had no motion cue, only the extracted density cue. With the density cue alone, only 1 of our 3 subjects achieved above-chance shape identification. On the other hand, with the motion cue alone (without the density artifact), shape identification was almost as good as with motion and density cues together. The slight impairment was probably due to the scintillation that accompanied removal of the density artifact.

Scintillation

The critics point out that our method of eliminating the density cue introduces a region-specific variation in scintillation, which is itself a possible cue in a shape identification test. They prescribe avoiding practice and eliminating feedback as the remedy. Are restricting the opportunity to practice and eliminating feedback the optimal methods of dealing with artifactual cues? Whatever the degree of practice or feedback, we believe that it is better either to determine the possible effectiveness of significant artifacts or to eliminate them, as we did the density cue. Simply assuming that lack of practice will suffice is not adequate. Indeed, the possibility that scintillation variations might be a shape cue was considered in the source article; it was dismissed because scintillation is an even weaker cue than the extremely weak density cue. That is, when displays are constructed without motion cues but with only a density cue or a scintillation cue, it obviously is harder to perceive areas where scintillation differs from the average than where density differs from the average (Lappin, Doner, & Kottas, 1980). This is because random error in dot density as a direct cue to 3D slope occurs only because of sampling error (limited number of dots) and the fineness of the 182×182 pixel grid, whereas random error in the relation of scintillation to a frame-to-frame change in 3D slope occurs because of the coarseness of the 10×10 grid of local areas within which density was kept constant. That is, because of the way stimuli were constructed, scintillation density was an objectively less reliable cue to shape than was dot density. Like the dot-density cue, the scintillation cue in our displays can be measured alone and it can be compensated or masked.

Displays were constructed that had a pure scintillation cue, without the changing-density or motion cues. The only subject who was able to perform above chance with the isolated density cue also viewed the new displays. With a pure scintillation cue, it was clear that his performance in a shape identification task would have been even lower than that with a density cue, although we did not feel it was worthwhile to run a formal experiment. Conversely, displays were constructed with normal density-controlled KDE cues, but with extraneous scintillation added uniformly throughout the display to mask the scintillation cue. Shape identification in these displays (with the scintillation cue rendered ineffective) appeared essentially equivalent to normal KDE displays. However, adding extraneous scintillation reduces the signal-to-noise in the stimulus, and more extended observations undoubtedly would reveal a slight impairment—not due to the loss of the scintillation cue but to the added scintillation. The bottom line for displays that are not scintillation-corrected is that the residual scintillation cue could be used to

make some extremely coarse discriminations (e.g., there probably is more scintillation on the left than on the right of the display) that may support above-chance shape identification for extremely sophisticated subjects; it is not a significant factor when normal motion cues are available.

General Procedures for Dealing with Artifactual Cues

The procedures used to measure and to eliminate the dot-density and scintillation cues illustrate general principles. If a particular artifact yields above-chance guessing, (a) measure the strength of the artifactual cue in isolation and (b) construct displays in which the cue is eliminated, masked, or rendered useless by irrelevant variation. Creating the cue in isolation is useful because bounds on the possible strength of the cue can then be determined. For example, dot density and scintillation were extremely weak cues. Eliminating artifactual cues in the displays of interest is an ideal solution, but is not always possible. Thus, dot-density cues could be eliminated, but this introduced scintillation cues that could not be eliminated but could be masked by adding still more scintillation. It was not necessary to use the third general method of dealing with unwanted cues—introducing irrelevant variation. For example, irrelevant variation is used to eliminate motion extent as an artifact in velocity estimation (McKee, Silverman, & Nakayama, 1986).

In our experience, it has never been necessary or preferable to deal with possible artifactual cues by using naive subjects without feedback and hoping that the subjects do not use the artifactual cues. To review our previous discussion: the problem is that, for optimal performance, subjects must also learn to optimally use the relevant cues, and this requires practice with feedback.

Summary and Conclusions

1. The extraction of 2D relative velocity is a basic substrate for deriving 3D structure from dynamic visual stimuli for both the true KDE or KDE-alternative computations.
2. The 53-shape lexicon for our identification task presents an ecologically valid test of shape recovery (KDE) for complex depth surfaces.
3. Practice in the 53-shape task serves to optimize identification performance; however, practice is not necessary to immediately perceive vivid KDE.
4. Properly conducted, experiments with feedback can measure the limits of human capacity; experiments without feedback measure the ability of subjects to generalize from their past experience to the experimental stimuli.
5. Excluding feedback in KDE experiments does not eliminate the possibility that artifactual cues may generate a correct response, it merely confuses the issue.
6. Scintillation is an insignificant cue in the 53-shape stimuli.
7. Deriving high-resolution 3D structures from 2D dynamic displays requires the first-order motion processing system. In moving-dot displays such as ours, the second-order motion system cannot be used to solve tasks that require simultaneous access to velocities at more than two locations. Therefore, to isolate the KDE performance supported by first-order motion

processing, it is desirable to use stimuli that require computation of velocities at more than two locations.

References

- Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, 2, 284-299.
- Braunstein, M., & Todd, J. (1990). On the distinction between artifacts and information. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 211-216.
- Chubb, C., & Sperling, G. (1988). Drift-balanced random stimuli: A general basis for studying non-Fourier motion perception. *Journal of the Optical Society of America A*, 5, 1986-2007.
- Chubb, C., & Sperling, G. (1989). Two motion perception mechanisms revealed through distance driven reversal of apparent motion. *Proceedings of the National Academy of Sciences USA*, 86, 2985-2989.
- Dosher, B. A., Landy, M. S., & Sperling, G. (1989). Kinetic depth effect and optic flow: 1. 3D shape from Fourier motion. *Vision Research*, 29, 1789-1813.
- Geisler, W. S. (1989). Sequential idea-observer analysis of visual discriminations. *Psychological Review*, 96, 267-314.
- Green, D. M., & Swets, J. A. (1974). *Signal detection theory and psychophysics*. New York: Krieger. (Original work published 1966)
- James, W. (1890). *Principles of psychology*. New York: Holt.
- Landy, M. S., Doshier, B. A., Sperling, G., & Perkins, M. E. (1988). The kinetic depth effect and optic flow: II. First- and second-order motion. *Mathematical Studies in Perception and Cognition* (Tech. Mem. No. 88-4). New York: New York University.
- Landy, M. S., Sperling, G., Perkins, M. E., & Doshier, B. A. (1987). Perception of complex shape from optic flow. *Journal of the Optical Society of America A*, 4, 108.
- Lappin, J. S., Doner, J. F., & Kottas, B. L. (1980). Minimal conditions for the visual detection of structure and motion in three dimensions. *Science*, 209, 717-719.
- 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.
- Parish, D. H., & Sperling, G. (1987). Object spatial frequency, not retinal spatial frequency, determines identification efficiency. *Investigative Ophthalmology and Visual Science, ARVO Supplement*, 28, 359.
- Sperling, G. (1989). Three stages and two systems of visual processing. *Spatial Vision*, 4, 183-207.
- Sperling, G., & Doshier, B. (1986). Strategy and optimization in human information processing. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and performance* (pp. 2-1-2-65). New York: Wiley.
- Sperling, G., Landy, M., Doshier, B., & Perkins, M. (1989). Kinetic depth effect and identification of shape. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 826-840.
- Todd, J. T. (1988). Perceived 3D structure from 2-frame apparent motion. *Investigative Ophthalmology and Visual Science, Supplement*, 29, 265.
- van Santen, J. P. H., & Sperling, G. (1984). Temporal covariance model of human motion perception. *Journal of the Optical Society of America A*, 1, 451-473.
- van Santen, J. P. H., & Sperling, G. (1985). Elaborated Reichardt detectors. *Journal of the Optical Society of America A*, 1, 300-321.
- Watson, A. B., & Ahumada, A. J., Jr. (1985). Model of human visual-motion sensing. *Journal of the Optical Society of America A*, 1, 322-342.
- Wundt, W. (1905). *Grundriss der Psychologie* [Foundations of psychology.]. Leipzig, Germany: Englemann.

Received August 28, 1989

Revision received October 12, 1989

Accepted October 13, 1989 ■

Publication Practices and Scientific Conduct

The recent disclosures of fraud in the conduct of research, reporting of research, or both in a number of scientific disciplines have prompted a widespread program of self-examination of publication practices and ethics.

The editor joins with APA in reminding authors of the principles of good publication practices and scientific conduct. Prospective authors are directed to the *Publication Manual of the American Psychological Association* (3rd ed.) and to the "Instructions to Authors" printed in this issue. The requirements of data availability, replicability, authorship credit, ethical treatment of subjects, and primary publication of data are important—they are meant to ensure responsible science and appropriate use of scarce and valuable resources.