Near stereothresholds measured with random-dot stereograms using phase disparities

Saumil S. Patel, Ph.D.,[a] Michael T. Ukwade, O.D., Ph.D.,[b] Harold E. Bedell, Ph.D.,[c] and Vanitha Sampath, B.Opt., M.S.[d]

Colleges of [a]Engineering and [a,c]Optometry, and [a,c]Center for Neuro-Engineering and Cognitive Science, University of Houston, Houston, Texas; [b]Dr. Ukwade is currently in private practice in Kemah, Texas; and [d]Smith–Kettlewell Eye Research Institute, San Francisco, California

Background: Clinically, stereothresholds for random-dot (RD) stimuli are measured at near with a typical resolution of 20- to 40-seconds arc. In this article, we describe a method by which stereothresholds are measured using RD stimuli on a conventional computer monitor with sub-picture-element spatial resolution.

Methods: The RD stimuli consisted of individual left and right eye images, viewed haploscopically from 50 cm through orthogonal polarizers. Cross and uncrossed horizontal disparities as small as 6-seconds arc were produced by introducing appropriate phase disparities within the individual spatial frequency components of the RD stimulus. The method of constant stimuli was used to determine the stereothresholds for 20 normal adult observers.

Results: The mean stereothreshold across the 20 observers was 24.1 ± 16.6-seconds arc, with an average trial-to-trial variability of ±23%.

Conclusion: Stereothresholds of a few-second arc can be measured accurately from a near distance for RD stimuli, using a conventional computer monitor. A clinical test based on this technique would allow the measurement of global stereothresholds with very high spatial resolution.

Key Words: Clinical testing, global stereopsis, stereoacuity

Patel SS, Ukwade MT, Bedell HE, et al. Near stereothresholds measured with random-dot stereograms using phase disparities. Optometry 2003;74:453-62.

Wheatstone[1] demonstrated that human observers perceive compelling stereoscopic depth when the two eyes are presented with separate images that contain small horizontal disparities between the relative positions of common features. These horizontal position disparities simulate the slightly dissimilar views that result from the lateral separation between the two eyes, when viewing a natural three-dimensional scene. The smallest retinal image disparity that observers can attribute reliably to a nearer or farther object is the stereothreshold. Although the stereothreshold may be measured using "local" targets that contain identifiable features in the images presented to each eye, it is generally preferable to assess the stereothreshold using "global" random-dot (RD) targets.[2] In contrast to local stereotargets, RD targets have no identifiable monocular features and consequently provide no monocular cues that can indicate the direction of disparity when an observer is asked to make depth judgments. On the other hand, some patients can report the perception of depth in local stereotargets, but have difficulty or are unable to perceive stereoscopic depth in RD targets. The difficulty experienced by some observers in perceiving depth in RD targets could be considered a potential drawback for the use of RD stereo tests. However, this difficulty could result in part from the absence of salient monocular features in most RD targets, which decreases the likelihood of accurate vergence and creates uncertainty about where perceived depth should appear. Here we describe a new method that allows the measurement of stereothresholds at a near distance using commercially available computer monitors. The stereotargets that we describe include monocular features to guide vergence and to direct the observer's attention, but provide no information about the direction of stereoscopic depth.

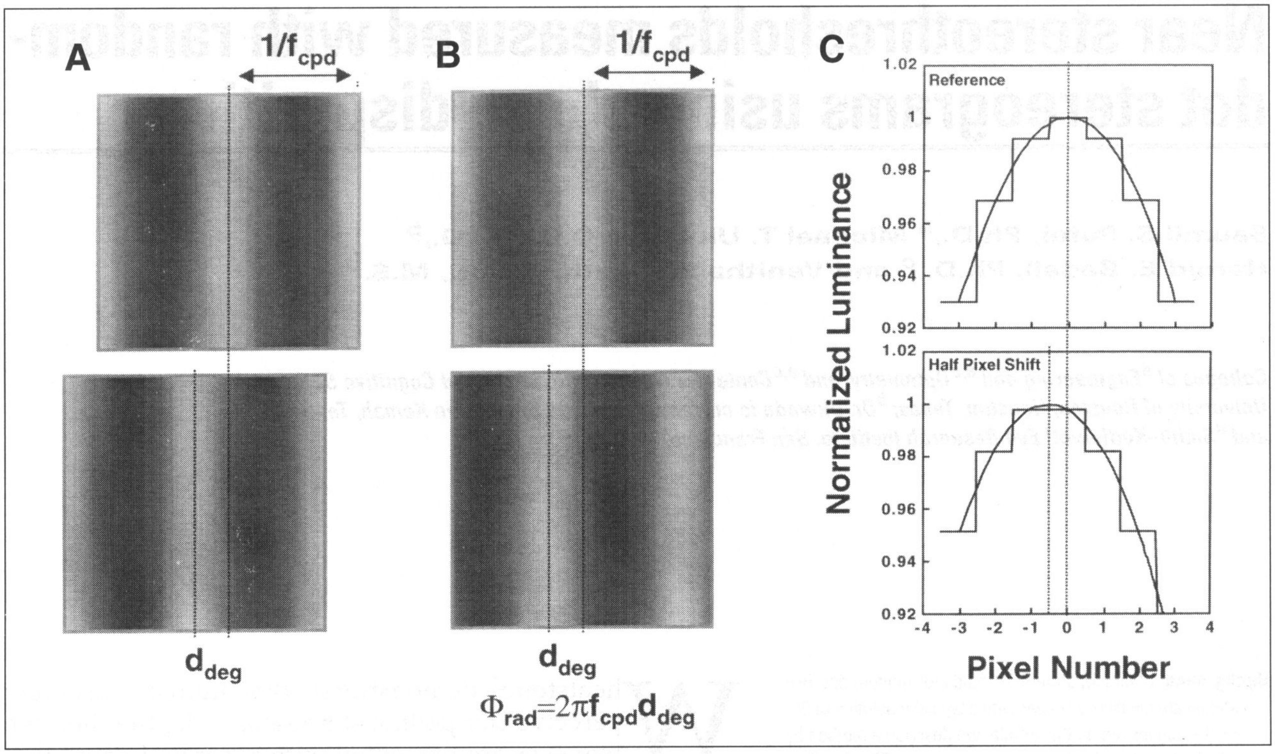

Figure 1 ***A***, Relative shift in luminance distribution introduced by shifting the position of one image relative to another. Each panel (top, bottom) shows a sinusoidal grating with a spatial frequency of f_{cpd}. Note that the subscripts for each of the variables in the figure indicate units. The relative difference in position of the two gratings is d_{deg}. The long dotted line that runs vertically through both gratings represents a horizontal reference position chosen arbitrarily. ***B***, Relative shift in luminance distributions introduced by shifting the phase of one distribution relative to the other. Each aperture contains two cycles of a sinusoidal grating of spatial frequency f_{cpd}. The grating in the bottom aperture is shifted leftward relative to that in the top aperture by d_{deg}. If we arbitrarily assign a phase angle of zero to the grating in the top aperture, then the phase difference Φrad between the gratings in the top and the bottom is given by $2\pi f_{cpd} d_{deg}$. When the image in the top aperture is seen by the left eye and the image in the bottom aperture is seen by the right eye, a crossed horizontal disparity of d_{deg} (equivalent to a "crossed" phase disparity of Φrad) will be produced. ***C***, Sub-pixel position shift of a grating stimulus. The top panel shows the normalized luminance of individual pixels around the long dotted line in the top grating shown in Panel B. The smooth curves represent the continuous luminance distributions that were sampled to obtain the luminance at each pixel number in the sampled luminance distribution (step-wise functions). The bottom panel shows the sampled luminance distribution corresponding to a half-pixel shift of the continuous luminance distribution in the top panel. Note that the separation between the two vertical dotted lines in this panel corresponds to a substantially smaller phase or position shift than is represented in Panels A and B.

In laboratory experiments, the optimal threshold for detection of relative retinal image disparity in RD stimuli is a few-seconds arc.[3-5] To measure these fine spatial thresholds in the laboratory, the binocular stimuli have to be presented on a display device with very high spatial resolution. To achieve the needed spatial resolution using available computer monitors, the stimuli are viewed typically from a considerable distance. However, because many daily tasks are performed at a near distance, and because some sensory and motor operations of the visual system are distance-dependent, measurement of the stereothreshold at a near distance is often more relevant for clinicians.

Previously, Stevenson, Cormack, and Schor[3] showed that *near* stereothresholds in the range of 3- to 5-seconds arc are possible using RD stimuli. These measurements were made using custom hardware that shifts the horizontal scan lines on a computer monitor by less than the size of a single picture element (pixel). The availability of such custom hardware is currently limited to research laboratories. Clinically, stereothresholds are measured at near using tests such as the Titmus, Randot, and Randot Preschool Stereoacuity tests, with a resolution of 20- to 40-seconds arc.[6,7] This resolution is inadequate to characterize the *optimal* performance of the stereovision system. The assessment of *optimal* stereothresholds may be helpful when clinicians wish to probe the functional consequences of subtle visuomotor dysfunctions and evaluate therapeutic interventions. For example, one consequence of correcting a sizeable fixation disparity at near might be to improve the stereoscopic depth threshold from 20- to 10-seconds arc,[8,9] which would be undetectable using most current stereotests. Similar considerations may arise when prescribing refractive correction, such as soft contact lenses, that may leave some residual blur. Other reasons to measure the optimal stereothresholds at near are to identify individuals who are faced with stringent visual demands. As examples, the performance of patients who are jewelers, tailors, or surgeons might be enhanced noticeably by an improvement of their near stereothresholds from 20-seconds arc to 10-seconds arc or better.

The resolution of a typical computer monitor is 72 dots-per-inch (dpi). This corresponds to a pixel size of about 140-seconds arc when the monitor is viewed from 50 cm. One technique for producing a horizontal disparity between a pair of monocular stimuli on a computer screen is to shift the identical luminance distributions that are presented to the two eyes in opposite directions, by an integer number of pixels (*see* Figure 1, *A*). However, a drawback of this technique is that the minimum horizontal disparity is limited to the size of one pixel. To measure stereothresholds at a near distance, one has to overcome this limited spatial resolution. One way to overcome this limitation is to use a single spatial frequency grating[10] or Gaussian blobs[11] as the binocular stimulus. As shown in Figure 1, *B*, the introduction of a horizontal position shift (d_{deg}) between a pair of grating targets is equivalent to introducing a phase shift (ϕ_{rad}) in one grating relative to the other. When the top grating is presented to the left eye and the bottom grating is presented to the right eye, the crossed horizontal disparity equal to d_{deg} is equivalent to a "crossed" phase disparity of ϕ_{rad}. To improve spatial resolution, this technique introduces a relative phase shift between the spatially sampled luminance distributions that are presented to the two eyes. Although the position of an image on a monitor can only vary by an integer number of pixels, the phase of a sinusoidally graded luminance distribution can vary virtually continuously. Figure 1, *C* illustrates how a continuous sinusoidal luminance distribution that is sampled at one-pixel intervals can be phase-shifted and then re-sampled to result in a position shift equal to one-half pixel. Extremely small spatial shifts can be introduced by re-sampling a sinusoidal luminance distribution, thereby reducing substantially the minimum horizontal disparity that can be presented between a pair of monocular images. Because this technique involves sampling, at fixed spatial locations, of sinusoidal gratings with a spatial phase difference between them, the images that are presented to the two eyes are not identical, except when the horizontal disparity is an integer multiple of the pixel size.

A disadvantage of using grating stimuli to measure stereothresholds is that—like line targets—gratings include strong local features that can provide unwanted monocular cues during binocular testing. More importantly, stereothresholds measured with gratings (as well as with Gaussian blobs and Gabor patches)[10-12] are an order of magnitude higher than the optimal threshold values that are measured with lines[13-16] or RD stereograms.[3-5] Indeed, Westheimer and McKee[17] reported that removing the information in any substantial band of spatial frequencies from a line stereogram resulted in a degradation of stereothresholds. Consequently, stereothresholds measured by introducing sub-pixel disparities between a pair of stimuli with limited spatial frequency content (such as gratings) are unlikely to reflect the optimal sensitivity of the stereovision system. Recently, Bach et al.[16] used sub-pixel disparities to measure stereothresholds of a few-seconds arc with a slightly blurred line stimulus.

Although it is relatively simple to introduce sub-pixel horizontal disparities between two grating stimuli, the process of introducing sub-pixel horizontal disparities in an RD stimulus is more complex.* Nevertheless, the elementary operations of this complex process are similar to those used for introducing sub-pixel horizontal disparities between two gratings. Like any repetitive or non-repetitive pattern, the half of an RD stereogram that is presented to one eye can be decomposed mathematically into a unique set of sinusoidal spatial frequency components. Equivalently, any repetitive or nonrepetitive patterned image can be produced by summing up a finite number of sinusoidal gratings, each of which is characterized by its spatial frequency, amplitude, relative position (or phase), and orientation. The method that is used most commonly to decompose a sampled image is called the *Discrete Fourier Transformation*. After this decomposition or transformation,

* Recently, Hess et al.[5] created sub-pixel disparities in stereograms created from random-spatial-noise images by interpolating between the luminance values of horizontally adjacent pixels. In some experimental conditions, this technique yielded stereothresholds in their two practiced observers that were less than 10-seconds arc.

each sine wave that contributes to the original image is termed a Fourier component. The Fourier transformation of a two-dimensional sampled image generates a two-dimensional Fourier representation that is specified mathematically as a matrix of complex numbers (A + Bi). This matrix of complex numbers contains the amplitude, phase, and orientation information for each Fourier component of the image. By reversing this process, a sampled image can be constructed from its Fourier components using a process called the *Inverse Fourier Transformation*.

To overcome the limited spatial resolution that is inherent in standard computer monitors, we presented sub-pixel horizontal disparities within pairs of RD stimuli by introducing the appropriate interocular phase disparities between the Fourier components of each eye's image. Small horizontal position disparities were represented by introducing these phase disparities in a spatial frequency- and orientation-dependent manner.

Methods

Observers

Twelve males and eight females voluntarily participated in this experiment. Most of the observers were optometry students or graduate students from the College of Optometry at the University of Houston. Informed consent was obtained before the experiments were conducted. In addition, our sample included four members of the college faculty and staff and two of their family members. The age of our subjects ranged from 8 to 50 years. All observers wore their habitual corrections and reported visual acuity of 20/20 or better at near and far in each eye. Each observer

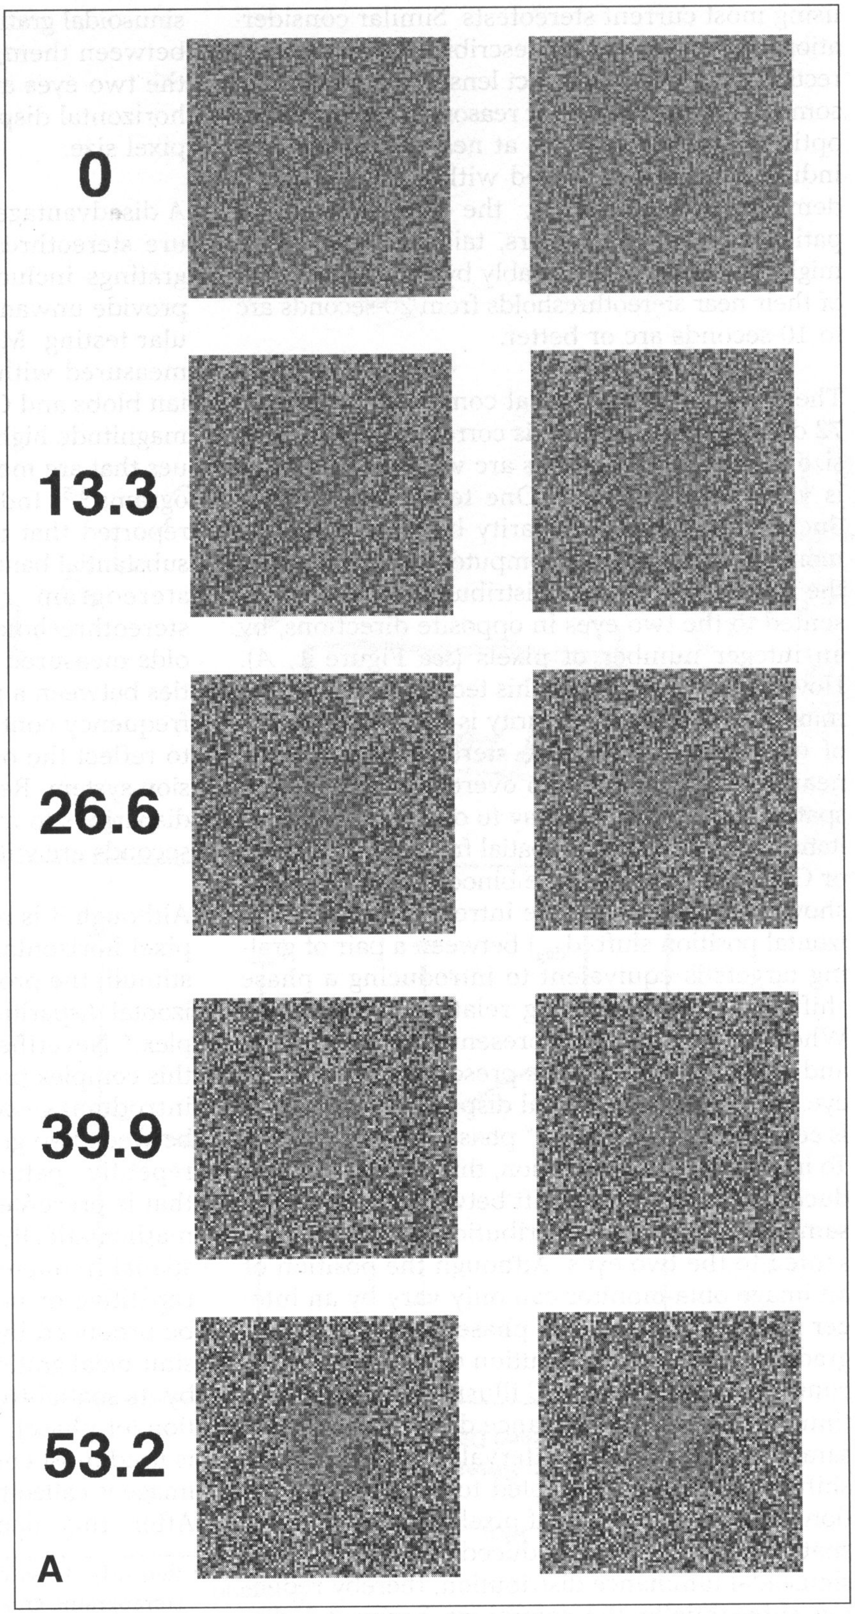

Figure 2 Sample pairs of images used in the experiments. If an image pair is fused with crossed eyes, the central square will appear ***A***, behind or ***B***, in front of the outer reference square. The numbers in the left column (signs are assigned arbitrarily) indicate the horizontal disparity within the central squares in seconds arc if the viewing distance is adjusted so that each outer square of RDs subtends a visual angle of 3.3 degrees.

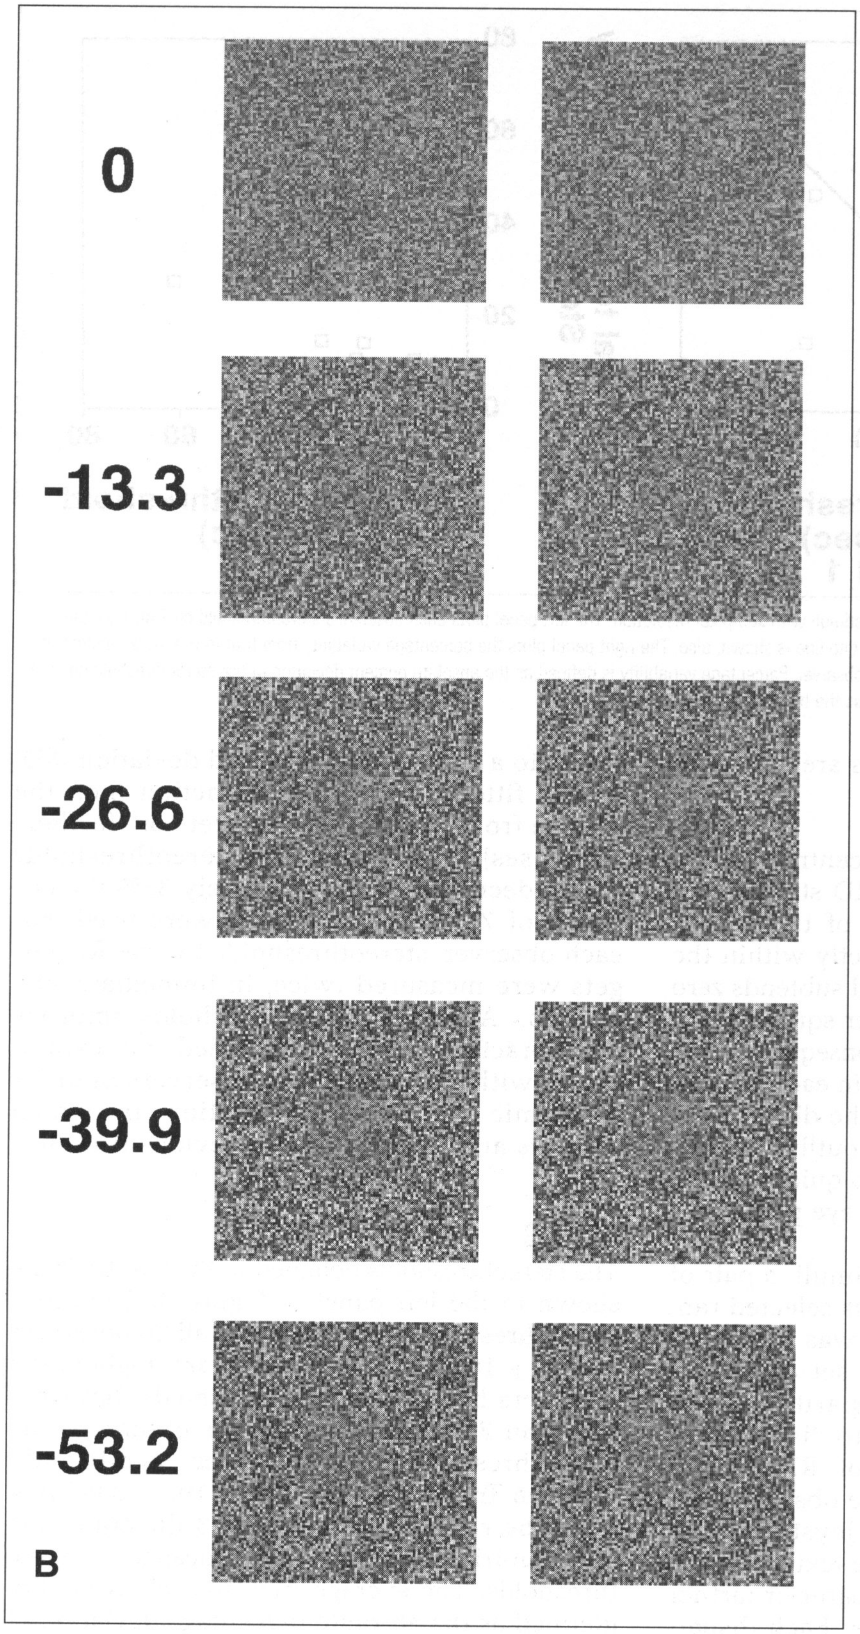

Stimuli

The RD stimulus consisted of pairs of individually constructed left and right eye images, viewed haploscopically in a dark room from 50 cm. The computer monitor was divided into two halves with each half covered by a polarizing film. The angle of polarization on the left half was 45 o[rientation]deg and that on the right half was -45 odeg. Observers wore a spectacle frame with -45 odeg and 45 odeg polarizing film in front of the left and the right eye, respectively. Because the two halves of the stereogram were separated laterally on the computer monitor and cross-fused, additional 4.5 p.d. base-in prisms were introduced in front of each eye to make the vergence demand approximately equivalent to a target viewed at a distance of 50 cm. The luminance of the targets presented on the computer monitor was calibrated using a Minolta LS-110 photometer. To minimize the effect of residual non-linearities after the luminance calibration, only the gray levels between 64 and 192 (out of the total range of 256 gray levels) were used to construct the stimuli. Each eye's image (50% contrast, mean luminance with polaroids in place = 7 cd/sq-m) consisted of a 1-degree central square of random dots embedded in a 3.3-degree square of random dots. Each dot subtended an angle of 2-minutes arc. For the experiments reported here, crossed and uncrossed horizontal disparities as small as 6-seconds arc were produced by shifting the phases of all spatial frequency components in the central square that was seen by the right eye with respect to the left eye. Additional details about how the RD stimuli were constructed are provided in the Appendix. Sample images of RD stimuli containing different

demonstrated stereoacuity of 100-seconds arc or better (average = 39 $\pm$ 27-seconds arc) when assessed with a laboratory-based test using line targets.[8]

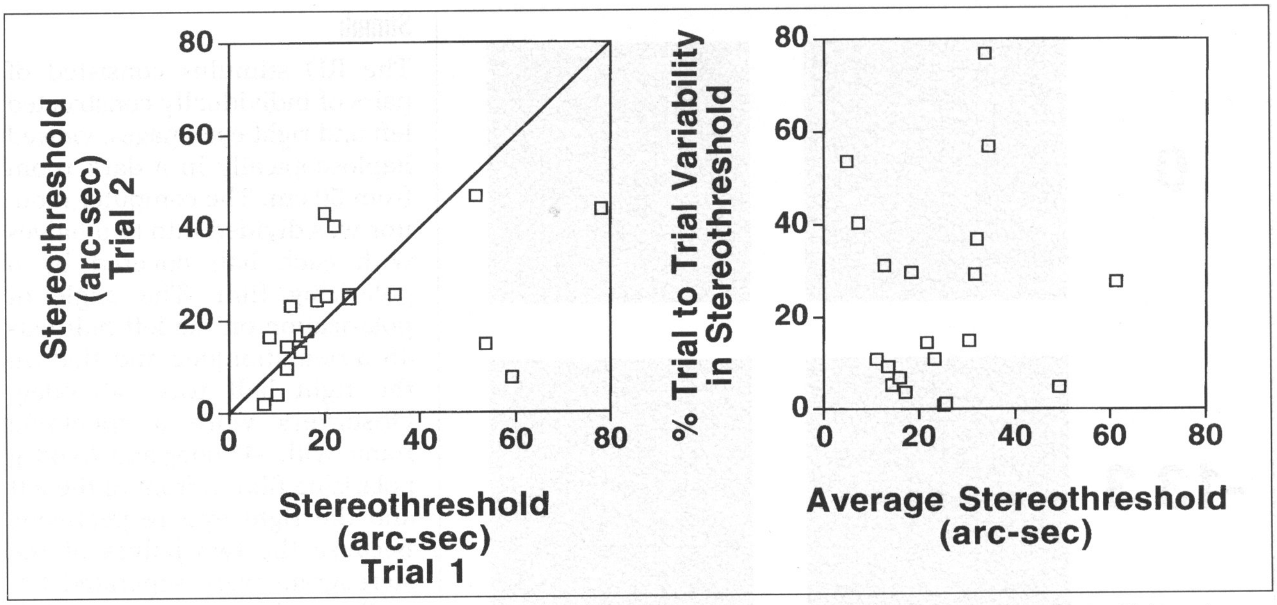

Figure 3 Stereothresholds measured using RD stimuli with sub-pixel resolution. The left panel plots each observer's stereothreshold on Trial 1 *vs.* the stereothreshold on Trial 2. The one-to-one line is shown, also. The right panel plots the percentage variability from trial to trial as a function of the average stereothreshold for each observer. Percentage variability is defined as the absolute percent deviation of the stereothreshold on Trial 1 from the mean stereothreshold across the two trials for each observer.

crossed and uncrossed disparities are shown in Figure 2, *A* and *B*.

Note that the outline of the central square remains visible in most of the RD stimuli presented in Figure 2. The outline of the central square, however, is centered exactly within the images presented to both eyes and subtends zero disparity with respect to the outer square when the two eyes' images are fused. Consequently, the central square that is visible within each monocular image provides no cue for the direction of perceived depth. However, the outline of this square may help the observers to quickly stabilize their vergence and versional eye positions.

Using the method of constant stimuli, a pair of RD images (one for each eye) was selected randomly from a set of images and was presented on the computer monitor. The set of images contained nine different disparities (four crossed, four uncrossed, and zero disparity) in the inner square. Each pair of RD images remained on the screen until the observer registered a response using a joystick. The observers indicated whether the texture in the central square appeared to be nearer or farther than the outer reference square. Each disparity was presented 10 times within one run, after which a psychometric function was fit to the observer's responses. Stereothresholds correspond to a change of 1 standard deviation (SD) on the fitted psychometric function (i.e., the change from 50% to 84% "nearer" or "farther" responses). The reported stereothresholds would decrease by approximately 33% if a criterion of 75% instead of 84% were used. For each observer, stereothresholds for the RD targets were measured twice, in immediate succession. Although stereothresholds improve with practice,[18,19] we performed our experiments without training the observers in order to mimic the clinical situation in which patients are tested without previous training.

Results

The stereothresholds obtained in the two trials are shown in the left panel of Figure 3. The mean stereothreshold averaged across all 20 observers is 24.1 ± 16.6 (± 1 SD) seconds arc. Eight of the observers had mean stereothresholds that were less than 20-seconds arc. Across observers, the stereothresholds do not differ significantly between Trial 1 and Trial 2 ($F[1,19] = 2.96$, $p = 0.1$). The right panel in Figure 3 illustrates the between-trial variability in the measured stereothresholds. The average trial- to-trial variability (defined as the absolute percentage deviation of the threshold on Trial 1 from the mean of the thresholds on Trials 1 and 2 for each observer) is ±23%.

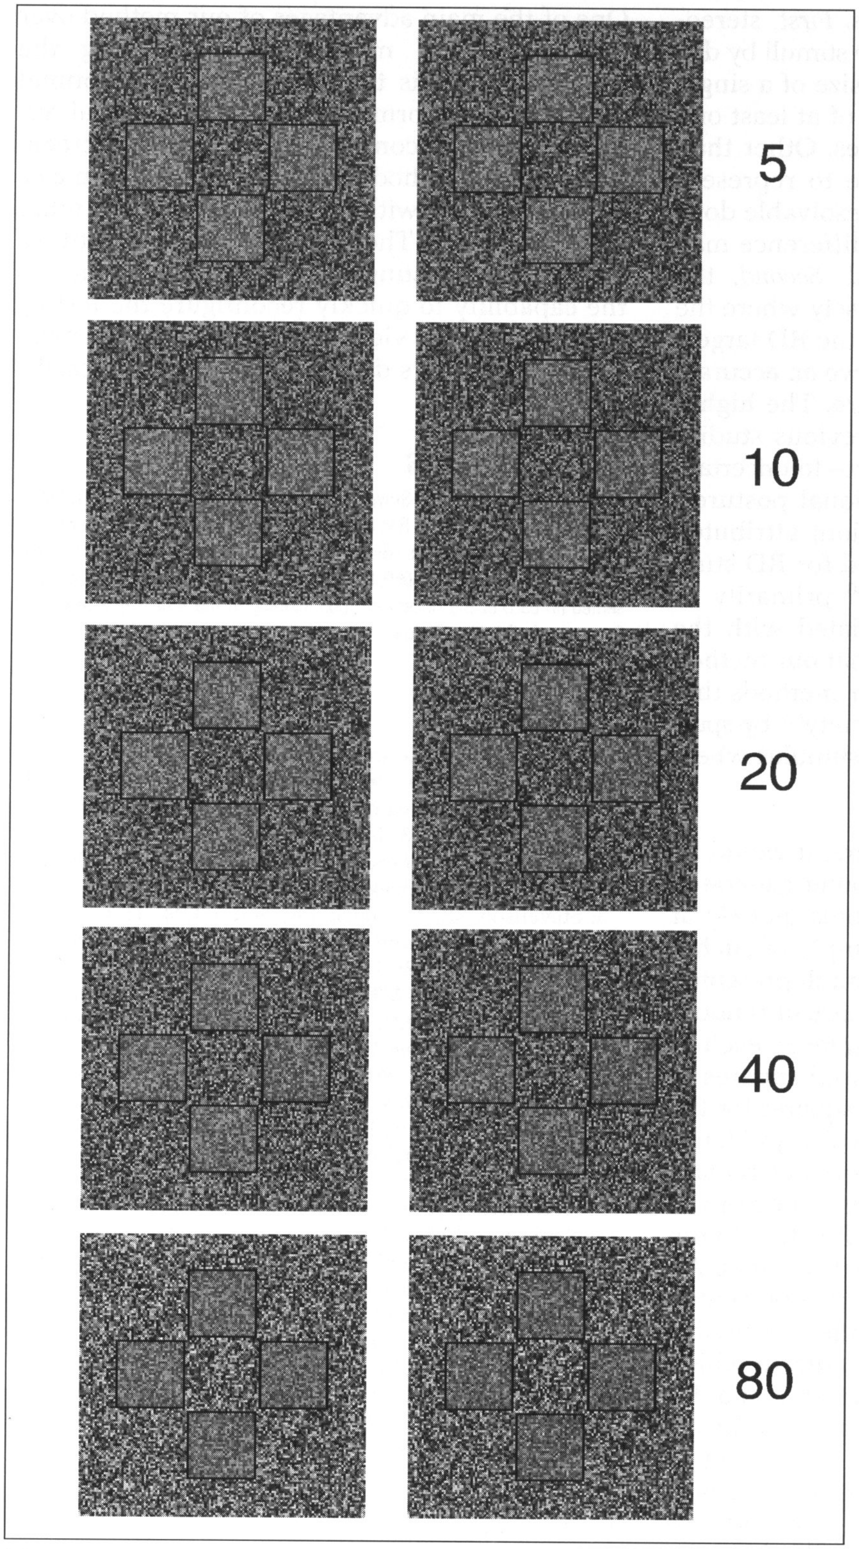

Figure 4 A sample chart to measure stereothreshold from a near distance. Upon fusing a pair of images on one row with crossed eyes, one of the four inner squares will appear in front of the other three inner squares. If viewed from the distance at which each of the four inner squares subtends 1 degree, the relative disparity between the inner square with crossed disparity and the remaining three inner squares with the same uncrossed disparity is shown on the right side. The correct answers, for the image pairs going from top to bottom, are the right, bottom, left, bottom, and top inner square.

Discussion

Stereothresholds in the hyperacuity range can be achieved at a near testing distance using RD stimuli created on a standard computer monitor with sub-pixel resolution. These fine thresholds were obtained despite the fact that 16 of the 20 observers were unfamiliar with the RD stereo task. The average trial-to-trial variability of both measurements was about 6-seconds arc. Trial-to-trial repeatability is particularly relevant if such measurements are to be made clinically. Although the stereothresholds measured on the first and second trials did not differ significantly, most of the observers' stereothresholds were higher on the first trial. For example, an improvement of stereothreshold was found on the second trial for 14 of the 20 observers, suggesting that a small amount of training may be necessary to achieve stable, low stereothresholds using these stimuli. In other words, a clinician may need more than one reading in order to obtain a precise measurement of a patient's baseline stereothreshold.

Some previous studies reported stereothresholds measured with RD stereograms that were substantially higher than the hyperacuity range. For example, Harwerth and Rawlings[20] found stereothresholds for RD stereograms that ranged from 15- to 200-seconds arc. In two additional studies, the mean stereoacuites for RD stimuli were reported to be 137 $\pm$ 87 (1 SD) seconds arc[21] and 40 $\pm$ 7.2 (1 SE) seconds arc.[22] In contrast, we found an average stereothreshold measured with RD stereograms of approximately 24-seconds arc. At least two differences between our study and the previous studies might contribute to the dissimilar outcomes. *First*, stereoscopic depth was created in our RD stimuli by disparities that are smaller than the size of a single dot, as compared with disparities of at least one random dot in the previous studies. Other than the fact that our stimuli are able to represent smaller disparities using readily resolvable dots, it is not clear how this stimulus difference may impact upon the stereothreshold. *Second*, the observers in our study knew precisely where the disparities would be presented in the RD targets and, therefore, could quickly achieve an accurate vergence and versional eye posture. The higher thresholds with RD targets in previous studies may be attributable—at least in part—to uncertain or inaccurate vergence and versional postures. Harwerth (personal communication) attributed the high stereothresholds obtained for RD stimuli in his study with Rawlings[20] primarily to methodological difficulties associated with the experimental set-up. We believe that our method gains a clear advantage over other methods that are used to measure global stereoacuity[6,7] by specifying explicitly the location in the stimulus where disparity will be presented.

Using the technique described here, it would be possible to create a series of computer-generated "charts" to measure stereothresholds quickly at near in a clinical setting. An example of such a chart is shown in Figure 4. The stimuli presented to each eye consist of an outer square of random dots with four embedded inner squares, each of which is similar to one of the central squares in Figure 2. As in Figure 2, the outer squares for the right and left eyes in Figure 4 are always identical. The outlines of all the inner squares remain at zero disparity, but the disparity in one of the four inner squares is opposite in direction to the disparities in the other three inner squares. An increasing magnitude of disparity is represented in the four inner squares of the successive "charts." The patient's task is to indicate which of the four inner squares on each chart is at a different depth (in front, for each of the five image pairs in Figure 4) than the other three, and to specify the location in depth (nearer or farther) of the different one. Because these judgments are between targets that lie in front of and behind the reference plane, it is logical to define the patient's stereothreshold as the disparity between the inner squares, which is equivalent to twice the disparity of each inner square from the reference plane.

One of the main advantages of our method over existing clinical methods of measuring the stereothreshold is that it can measure optimal stereovision performance at very high spatial resolution using a conventional computer screen. Because the method is computerized, it can easily be integrated with computerized tests of other visual functions. There are numerous benefits of computerized testing of visual functions, such as the capability to quickly reconfigure the testing parameters (e.g., viewing distance), and to manage each patient's data more quickly and easily.

Acknowledgments

We wish to thank Dr. Scott Stevenson for his help in constructing the RD stimuli used in our study and Dr. Ron Harwerth for helpful discussions. We also wish to thank all of the observers who generously volunteered their time. This work was supported by R01 EY05068, R01 MH49892, and P30 EY07551, and was presented (in part) at the December 2000 meeting of the American Academy of Optometry in Orlando, Florida.

References

1. Wheatstone C. Contributions to physiology of vision — Part the first. On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Phil Trans R. Soc London* 1838;B2:371-93.
2. Julesz B. *Foundations of Cyclopean perception*. Chicago: University of Chicago Press, 1971.
3. Stevenson SB, Cormack LK, Schor CM. Hyperacuity, superresolution, and gap resolution in human stereopsis. *Vision Res* 1989;29:1597-605.
4. Patel SS, Ukwade MT, Bedell HE, et al. Near stereothresholds measured with sub-pixel resolution using random-dot stimuli. *Optom Vis Sci* 2000;77(Suppl):270.
5. Hess RF, Liu CH, Wang Y. Luminance spatial scale and local stereo-sensitivity. *Vision Res* 2002;42:331-42.
6. Eskridge JB, Amos J, Bartlett JD. *Clinical procedures in optometry*. Philadelphia: Lippincott, 1991:124-30.
7. Birch E, Williams C, Hunter J, et al. Random dot stereoacuity of preschool children. *J Pediatr Ophthalmol Strabismus* 1997;34:217-22.
8. Ukwade MT, Bedell HE, Harwerth RS. Stereopsis is perturbed by vergence error. *Vision Res* 2003;43:181-93.
9. Kromeier M, Schmitt C, Bach M, et al. Stereoacuity versus fixation disparity as indicators for vergence accuracy under prismatic stress. *Ophthalmic Physiol Opt* 2003;23:43-9.
10. Morgan MJ and Castet E. The aperture problem in stereopsis. *Vision Res* 1997;19:2737-44.
11. Schor CM and Wood I. Disparity range for local stereopsis as a function of luminance spatial frequency. *Vision Res* 1983;23:1649-54.
12. Westheimer G. Lines and Gabor functions compared as spatial visual stimuli. *Vision Res* 1998;38:487-91.
13. Berry RN. Quantitative relations among Vernier, real depth, and stereoscopic depth acuities. *J Exp Psychol* 1948;38:708-21.
14. Stigmar G. Observations of Vernier and stereoacuity with special reference to their relationship. *Acta Ophthalmol* 1970;48:979-98.

15. McKee SP. The spatial requirement for fine stereoacuity. *Vision Res* 1983;23:191-8.
16. Bach M, Schmitt C, Kromeier M, et al. The Freiburg stereoacuity test: automatic measurement of stereo threshold. *Graefes Arch Clin Exp Ophthalmol* 2001;239:562-6.
17. Westheimer G and McKee SM. Stereoscopic acuity with defocused and spatially filtered retinal images. *J Opt Soc Am* 1980;70:772-8.
18. Wittenberg S, Brock FW, Folsom WC. Effect of training of stereoscopic acuity. *Am J Optom Arch Am Acad Optom* 1969;46:645-53.
19. Fendick M, Westheimer G. Effects of practice and the separation of test targets on foveal and peripheral stereoacuity. *Vision Res* 1983;23:145-50.
20. Harwerth RS and Rawlings SC. Viewing time and stereoscopic thresholds with random-dot stereograms. *Am J Optom Physiol Opt* 1977;54:452-7.
21. Mehta AM and France TD. Distance stereoacuity in children and teenagers with normal near acuity. *Am Orthop J* 1997;47:144-7.
22. Schmidt PP and Brown KL. Comparisons of stereothreshold for contour and random dot targets. *Optom Vis Sci* 1995;72(Suppl):85.

Corresponding author:

Harold E. Bedell, Ph.D.
University of Houston
College of Optometry
505 J. Davis Armistead Building
Houston, Texas 77204-2020

HBedell@Optometry.UH.edu

Appendix

Construction of random-dot stimuli

Here we describe how the central square in the right eye's image was constructed with a specific disparity from the central square of the left eye's image (*see* Figure 2). First, the RD image that comprised the central square in the left eye's image was transformed mathematically to its Fourier representation. This representation was a matrix that specified the amplitude and phase of each sinusoidal spatial-frequency component (0 to 21 cpd) of the left eye's central square, at each orientation between –90 and 90 degrees of orientation. Our angular convention is that a vertical sine wave grating is at 0 degree of orientation. Horizontal disparity was introduced into the right eye's image by adding a spatial-frequency and orientation-dependent phase shift to each spectral component of this Fourier matrix. The phase shift (ϕ, in radians) that was added to the spatial frequency component (f, in cpd) oriented at an angle (α) with respect to the vertical is given by

$$\phi(f, \alpha) = 2\pi df \cos(\alpha)$$

in which, d is the desired horizontal disparity in degrees. This equation was applied to all positive spatial frequencies (i.e., the elements of the Fourier matrix in Cartesian quadrants 1 and 4). To ensure that the inverse transform of the phase-shifted Fourier matrix remained real, mirror symmetric elements on the negative spatial frequency side of the matrix were set to be the complex conjugates (i.e., the imaginary part changes sign) of the elements on the positive frequency side. The inverse Fourier transform of this manipulated Fourier matrix yielded the central square of the right eye's RD image. Both eyes' inner squares had the same amplitude spectrum and, consequently, an identical mean luminance of 7 cd/sq-m in the central regions. The mean luminance of the outer region in both eyes' images is equal to the average of the minimum and maximum luminances present in both of the inner squares. This contrast scaling introduces a slight difference in contrast between the inner squares and the outer region in both eye's images, so that the inner squares in the left and right eyes have a similar appearance with respect to the outer squares. All of the RD images used in our experiments were constructed using MatLab software (Mathworks, MA). The smallest horizontal disparity that we produced between the RD targets in our experiment was 6-seconds arc. The smallest disparity that can be produced with our technique is considerably smaller than a pixel and is limited only by the dynamic range of contrast that can be produced on the computer monitor.