

May 01, 2010

Seeing many depth planes simultaneously

David A. Lynch
Northeastern University

Recommended Citation

Lynch, David A., "Seeing many depth planes simultaneously" (2010). *Honors Junior/Senior Projects*. Paper 59. <http://hdl.handle.net/2047/d20000670>

This work is available open access, hosted by Northeastern University.

Seeing Many Depth Planes Simultaneously

David A. Lynch

Dr. Adam Reeves

Department of Psychology, Northeastern University, 125 NI

Honors Thesis: Reeves' Lab (239 LA)

4/27/2010

Abstract

Depth perception allows us to navigate a three dimensional world. To enable depth perception, our visual system uses several depth cues, one being stereopsis. *Stereopsis* results from the fact that humans have two forward-facing eyes that receive different views, which are then combined into one cohesive image with depth. By programming displays with multiple disparities, multiple transparent depth surfaces can be created with each disparity applying to a fixed number of stimulus elements (Random Dot Stereograms; RDS). Transparent refers to more than one depth surface in the same local area of the visual field, like leaves on a tree. Using RDS displays, Wilcox et al (2008) has shown that up to 6 depth planes can be reported almost perfectly with unlimited viewing time. However, unlimited viewing time permits multiple eye movements, so an impression of depth can be built up over time even if only 2 or 3 depth planes can be seen at once. We therefore re-ran her experiment with viewing time limited to 400 ms, a display time that made the display appear and disappear before an eye movement could be made, to see how many depth planes can be seen simultaneously (Semlow & Wetzell, 1979; Satguman et al, 2009). We also asked if people could see depth from opposite-contrast images (i.e. black dots corresponding with white dots). According to Harris & Parker (1995), this is not possible. Supplementing previous research (Wilcox et al, 2008), the findings demonstrated that the human visual system, under optimal conditions and after learning, is capable of perceiving 6 or more depth planes simultaneously. Contrary to the findings by Harris et al (1995), the human visual system is capable of acquiring depth information (up to 6 depth planes) within stimuli with same and opposite contrasts. We have found that the learning necessary to successfully distinguish these novel depth displays in adults is slow and progressive, possibly resembling the learning of depth cues in infancy.

The phenomenon of depth perception allows us to navigate a complex three dimensional world with our eyes. To enable depth perception, our visual systems must decipher several depth cues, as first theorized by Berkeley (1709). One such depth cue is stereopsis, discovered by Charles Wheatstone (1838). Stereopsis results from the fact that humans have two forward-facing eyes that receive unique input, which is then combined into a cohesive visual experience. The differences between the images in the eyes results in *disparity*; such differences are contingent on the viewing angle and the depth of the surface. By increasing disparity, up to a value known as “Panum’s limit,” the sensation of depth will be increased. Beyond Panum’s Limit, or by crossing our eyes, we experience “double vision” by forcing corresponding point to separate. As an example of pure stereoscopic perception, consider an instance devoid of depth information from overlapping objects, shadows, perspective, and other pictorial depth cues. Imagine someone throws a Frisbee in such a way that it hovers far above your head before descending towards the group. You look up into a vast blue sky. Without depth clues that allow the comparison between relative distances, you rely solely on binocular stereopsis to judge the Frisbee’s distance. Figure 1 shows a schematic of how the input from each eye is integrated into a single image.

Ogle (1950) summarized a century of research in stereopsis which was based on presenting recognizable stimuli, such as pictures, to the two eyes. He concluded that the brain integrates representations of the stimuli, one from each eye, in order to compute the depth relations between the objects in the scene. This conclusion was shown to be false by Dr. Béla Julesz, who in 1959, pioneered the technique of the Random Dot Stereogram (RDS). By using displays constructed from randomly-located dots, Julesz was able to portray stereoscopic figures in depth whose monocular images were completely

unrecognizable – just randomly placed dots. Nevertheless, the stereoscopic figures were easily recognizable. Therefore the dots themselves, rather than higher-level representations, had to be integrated across the two eyes. In his groundbreaking book *Foundations of Cyclopean Perception*, Julesz (1973) described the phenomenon of the human visual system utilizing the input of two eyes to perceive one coherent image without requiring prior recognition.

In the classic Julesz RDS displays, the disparities of the dots were selected so that different surfaces would be seen at different depths. In the simplest of such displays, all the dots in a central square region of the display were given a different disparity than all the remaining dots in the rest of the display. In this case, the observer will see the central region as a surface, either in front or behind the remaining dots which form a second surface, just as if one were looking down at the top surface of a match-box placed on top of a desk. Julesz discovered that more complex surfaces could also be portrayed with random dots. Observers can see tilted surfaces extending in depth away from them, including a spiral figure (Julesz, 1973), not just several flat surfaces.

These facts suggest that stereopsis is computed ‘locally’ in the brain, that is, each small region of the visual field gives rise to its own local average disparity which determines its own relative depth. If so, it should be impossible to see more than one depth plane in the same local region. Yet, if one looks at leaves on trees, one has the distinct impression of seeing both the leaves on the tree and some of the landscape beyond, all in the same region of visual space. That is, one can see ‘through’ the leaves to what is behind, and be aware of both the leaves and what is behind them at the same time. One explanation for this is that one can accommodate the lenses of the eye, and alter the vergence angle between the eyes, to emphasize visual detail in one depth plane and blur

away detail at other depth planes; thus one could in principle verge on the leaves and then, later on, verge of the landscape behind. However, Julesz's RDS method allows the experimenter to control for such optical effects, and it is nevertheless the case that one can see a set of random dots with disparities defining a distant plane through another set of random dots with disparities defining a forward plane, even when all the dots are spatially intermingled rather than grouped into separate surfaces (Zabulis & Backus, 2004). This surprising result motivated us to ask, just how many depth planes can be seen, if the dots defining the planes are spatially intermingled?

It is possible in the RDS displays to program intermingled dots to portray many depth planes at once, using disparity increments to "pile on" the depth planes. The depth planes now appear transparent. The percept resembles looking through several smudged window panes stacked along one's line of sight from proximal to distal in relation to the viewer. Each piece of glass represents a single transparent depth plane, and the smudge represents a single point in the depth space. By programming RDS displays with multiple disparities, multiple transparent surfaces can be created with each disparity applying to a fixed number of stimulus elements.

Wilcox et al (2008) were interested in the number of transparent RDS depth planes the visual system can perceive. Wilcox's data showed that participants may perceive up to 6 such depth planes (Wilcox et al, 2008). However, her data did not demonstrate that the observers saw all the depth planes simultaneously. Since she presented RDS to participants for as long as they wanted, typically 3-5 seconds, several vergence eye movements may have been made, which would allow the observer to parse the display in a piecemeal fashion. Vergence eye movements, a particular type of eye movement, change the viewing angles of the retinae in order to line up the corresponding

points in space. Our eyes continually make movements, or saccades, to align the desired viewing area within the retina. Since gross movements of the eye result in de-focusing, each saccade must be quick enough so it does not hinder the visual experience. We experience vergence when our pupils turn inward when we examine an object close to our face (converge) and our pupils move away from one another when viewing objects in the distance (diverge). When applied to the Wilcox (2008) experiment, each eye movement could have resulted in a recalculation of depth information. By shifting the eyes, a participant may perceive, say, just 2 or 3 depth planes, while the remaining dots forming 'depthless' visual noise. Then the observer may re-verge their eyes to investigate the noise, subsequently discovering 2 or 3 more depth planes, and so on. This participant can report 6 depth planes; however, the visual system only parsed the display in multiple steps. Although the study provided evidence that 6 depth planes may be accurately reported, only a maximum of 3 depth planes were seen simultaneously. Therefore, her results can not address the question of how many depth planes one can see all at once.

In our study, we prevented vergence eye movements to control for this possibility. Because vergence eye movements take time to execute, we were able to present RDS displays that would appear and disappear before an eye movement could be made. Vergence movements start about 160-200 ms after onset of a display (Semlow & Wetzell, 1979), but are relatively slow to execute, reaching peak velocity only at about 372 ms after onset (Satguman et al, 2009). Thus, by limiting the display time to 400 ms, no vergence eye movement was likely to have been completed. Since a change of vergence would temporarily disrupt correspondence, it is not advantageous for our observers to use more than one fixation when determining the number of planes in the display unless the

display remains on for a considerable time. Based on the time it takes to execute an eye movement, the participants in the Wilcox (2008) experiment with a 5 sec display time may have made up to 10 eye movements. In our experiment, we wanted to prevent vergence eye movements (by using a 400 ms display time), and thus examine the number of depth planes that may be perceived simultaneously.

We typically examined RDS displays with the same contrast dots in each eye; for example, both eyes view white dots on a black background. However, we also examined the perception of RDS displays with opposite contrasts, for example, with white crosses in the presence of potential interference in the form of randomly-located black crosses, all presented on a black background. According to Harris & Parker (1995), it is not possible to discern cohesive depth information from opposing-contrast stimuli (i.e. black and white dots that form corresponding points). Several researchers contend that our visual system evolved not to bind erroneous corresponding points (e.g., Smallman & McKee, 1995). According to this theory, the opposing contrasts do not register as ‘like terms’ in the cortex and subsequently they will be discarded. If the visual system parses light increments and light decrements through separate ON and OFF pathways, as reported by physiologists studying visual cortex of monkeys, then additional black crosses should have little or no effect on the perception of RDS stimuli composed of white crosses. However, displays in which each eye receives opposite-contrast crosses should not generate depth percepts at all. An alternative hypothesis relies on the view that depth perception requires learning (Berkeley, 1709). Since natural viewing only rarely provides each eye with opposite contrast images, the visual brain presumably has not had the experience needed to analyze such information. We therefore provided extensive

training on opposite contrast images. Our study may provide insight as to whether training facilitated a rarely-used spatial ability.

Since many people were not accustomed to RDS displays, particularly with a brief 400 ms display duration, training was necessary even with same-contrast images. Visual learning became an important aspect of the experiment to implement and track. Over the course of months, participants gradually reduced their display time as their accuracy increased. For ‘learned’ participants, The 400 ms display time results were initially the data of interest, but the learning that appeared to occur became an important ancillary finding. In this study we set out to determine the number of depth planes that can be seen simultaneously by preventing vergence eye movements.

I. Problem: Summary

As already stated, Wilcox et al (2008) has shown that up to 6 depth planes can be reported almost perfectly with unlimited viewing time. However, unlimited viewing time permits multiple eye movements, so an impression of depth can be built up over time even if only 2 or 3 depth planes can be seen at once. We therefore re-ran her experiment with viewing time limited to 400 ms so that multiple eye movements could not be made during viewing. In this study, we wanted to build on Wilcox’s findings and examine how many depth planes the visual system can parse simultaneously.

We also asked if people could see depth from opposite-contrast images. According to Harris & Parker (1995), this is impossible. Thus, the current theory assumes that depth comes only from points in the left and right eyes with the same contrasts (i.e. both white or both black). We ran participants with stable learning on RDS stimuli with opposite-contrast information given to each eye to see if the currently theory holds. If this contention proves to be valid for the cortical mechanisms processing depth

from stereopsis, then black crosses added to a display of white crosses should have little or no effect on the perception of depth. Using similar logic, if each eye receives opposite-contrast crosses, no depth should be seen at all.

II. Project Methodology and Resources

a. Participant Requirements

The participant pool, students between the ages of 18-22, was scheduled via the Northeastern University's PsyLink website. Each participant, prior to testing, successfully completed a standard eye exam demonstrating normal visual acuity, and perceived two depth planes when peering into a Julesz stereopsis apparatus. Glasses were permitted for the study. Participants who failed the eye chart were not penalized, but signed up for another experiment under PsyLink. All experiments were approved by the Northeastern University's Internal Review Board.

Across subject variability was one of the primary challenges of the study. Since visual learning is largely implicit, the experimenter as well as the subject has little indication as to how well the participant would perform. Each subject arrived at the study with varied depth experience. With the advent of the *Magic Eye* books and various other depth-related activities, previous 'training' is difficult to account for. The participants were, however, naïve to the RDS stimuli that we present.

b. Apparatus

Participants positioned their heads on a chin rest approximately 60 centimeters from the ViewSonic Professional Series P220f computer monitor. Adjustable 1.75 diopter lenses were set up to provide viewing of the stimulus display, based on the Brewster Method. The Brewster Method made the display surface so that it was seen at optical infinity; this made all potential depth planes equally visible, even though stimulus

elements are portrayed at a fixed distance (that of the TV monitor screen). Disparate dots were portrayed on the left and right-hand sides of the monitor. A vertical sheet of black cardboard was placed between the nose and the middle of the screen so that each eye saw its own half-screen. A 15 diopter prism lens was placed in front of the left eye to overlap the left and right visual fields so that disparate elements fused and depth became apparent. Adjusting the lenses and prism took approximately 10 minutes per participant since the optics required careful alignment with the optical axis of each eye. The participant could remove their head from the apparatus at any time and move back into the chin-rest without disturbing the set-up, so there was no reason for discomfort. Refer to Figure 1.

c. Program

Using custom code, we created RDS stimuli on the monitor via Windows-based MATLAB 6.5 and Cambridge Vision System hardware. With the program, we could adjust multiple variables in the program, including: the disparities, the number of dots per row, the dot contrasts, and the display duration. The computer monitor was programmed to display two random stimuli, each comprised of many small crosses, one stimulus in each visual field. The disparities between the two stimuli create the illusion of depth. The program allowed both fixed and randomized trials of RDS displays.

The standard program configurations conformed to Wilcox and Weishall's RDS displays (Wilcox et al, 2008; Weishall et al, 1989). Our displays consisted of: 12 dots per row, 3 pixels of disparity per additional depth plane, 8 columns, and variable display durations. The program was altered during later experiments to examine whether the visual system can parse opposing contrast displays. Opposing contrast displays were

programmed to have white dots in the left visual field, and corresponding black dots in the right visual field.

d. Methodology

Guided by previous findings (Wilcox, 2008; Weishall et al, 1989), preliminary experiments allowed us to evaluate the standardized parameters, which included disparity increment, dot density, and total number of dots per row. These findings led to our standardized experimental design. Once the parameters were finalized, we ran naïve participants through the RDS displays. Beginning with fixed displays, participants familiarized themselves with the stimuli before graduating to the randomized program. As learning progressed, the display duration was shortened.

Based on the previously mentioned observations by Wilcox et al (2008), our study utilized RDS stimuli to explore the identification and perception of transparent depth planes. Wilcox had presented RDS to participants for as long as they wanted, typically 3-5 seconds, which is enough time for several vergence eye movements. In order to prevent these eye movements, and thus discern how many depth planes may be perceived by the visual system simultaneously, we eventually (i.e. after much practice) utilized a viewing time of 400 ms, faster than the time needed to re-verge the eyes.

Since the ultimate test stimulus was so brief, only 400 ms, we used a priming display to provide the participant with enough time to verge his or her eyes properly. We presented a 2 second, prime display that preceded each 400 ms trial. The prime display contained just 2 depth planes, with the test display appearing immediately afterwards. If a participant missed the test display entirely or could not induce the sensation of depth, they were told to enter '0' and continue to the next display. The 'missed' display was removed from the data set. After practice, very few test displays were missed.

Test displays contained from 2 to 4 depth planes during initial practice, and then the maximum number of depth planes was slowly increased such that the observer would not make more than 10% errors at any stage. Initially, the number of depth planes was fixed for a block of trials and the observer knew what to expect, so that he or she could learn what multiple depth planes looked like. After this phase, the number of depth planes was randomized on each trial, so that the observer would not know what to expect. The randomization was controlled by an algorithm that equated the numbers of depth planes. Thus, for example, when 6 was the maximum number of depth planes, 2, 3, 4, 5, or 6 depth planes were equally-often presented.

Once learning had roughly stabilized, considerable data were obtained with 2 to 6 depth planes and later, with 2 to 8 depth planes. Some data were also collected with from 2 to 10 depth planes for the most experienced observers.

e. Statistical Analysis

After each trial, data were analyzed by regressing the number of depth planes reported against the number of actual depth planes presented. If all responses matched the number of depth planes presented, a correlation of 1, with a linear regression line of slope=1, would result. However, accurate reporting of low number of depth planes was sometimes accompanied by underestimation of high number of depth planes. In these cases, we utilized a cubic regression line to best-fit the data.

To track learning over time, we measure the amount of *bits lost* per trial block and then compared the values over time. Bits allowed for partial credit when erroneous responses were close to the actual stimulus (i.e. responding 7 depth planes given a stimulus of 8 planes). We added the total *bits lost* in each block of trials.

$$\text{Bits Lost} = -\sum [(\log_2 S) - (\log_2 R)],$$

Where S = the number of depth planes in the display and R = the number of depth planes reported.

The following table shows illustrative graded penalty:

S	R	Bits Lost
2	2	0.00
3	3	0.00
3	2	-0.58
3	4	-0.42
3	5	-0.74
3	6	-1.00
4	4	0.00
4	3	-0.42
4	2	-1.00
4	5	-0.32
5	5	0.00
5	4	-0.32
5	3	-0.74
5	2	-1.00
5	6	-0.26
6	6	0.00
6	5	-0.26
6	4	-0.58

III. Early Learning and fixed depth planes: Experiment 1

a. Problem

Using naïve subjects from the Northeastern University Undergraduate participant pool, we wanted to examine whether an untrained visual system could parse the transparent depth displays. The apparatus and program details were identical to the described set up in the “apparatus” section. Most participants experienced the displays during an hour long session; however, a couple participants experienced a 2 hour session.

b. Procedure

Beginning with fixed numbers of depth planes, participants began by previewing different numbers of depth planes with no time limit on their viewing. The research

would first present a 2 depth plane display and ask the participant to enter into the computer the number of depth planes seen. For the remainder of the block, the participant saw 2 depth plane displays one after another. Once the participant feels comfortable with 2 depth plane displays, they graduated to a 3 depth plane display block. Incrementally, the researcher increased the fixed depth planes by 1 as the participant became more comfortable with the displays. As mentioned prior, the number of depth planes remained fixed during these blocks of trials; however, the dots were randomized on each display. After the fixed depth plane trials were completed up to 8 depth planes, we began ‘previews’ of the time-limited randomization program.

Results

Since each participant was aware of the number of presented depth planes, they could search the display until their perceptual experience matched the number of depth planes presented. All participants reported being able to parse up to 8 transparent depth planes with various amounts of viewing time.

IV. Middle Learning and randomized program: Experiment 2

a. Problem

We wanted to examine whether the “learning” that occurred during the fixed depth plane displays transferred to the randomization program. The apparatus was identical to the described set up in the “apparatus” section; however, modifications were made to the programmed displays.

b. Procedure

Once participants reported that they felt ‘confident’ or ‘comfortable’ with the transparent depth displays, they were began on the randomized paradigm. The randomization program displayed stimuli with 2, 3, 4, 5, 6, 7, and 8 depth planes in a

random fashion, although each value was programmed to appear 5 times during the entirety of the block of trials. Beginning that random program was contingent on the participant's comfort level with the previously explained fixed depth plane displays. Each participant was given a 'trial run' of the paradigm to get them accustomed to the input keyboard and the priming stimulus. The participant entered the number of perceived depth planes into the computer using the before mentioned keyboard. Depending on the participants' response to the 'trial run,' we began all subjects with initial display times between 5 and 10 seconds. After multiple sessions, viewing time was gradually lowered 400 ms. Since the test stimulus was then so brief, we had to use a 2 sec priming display to provide the participant with enough time to verge his or her eyes properly before each trial, as explained in the Procedure section.

After each block of trials, Pearson's Correlation Coefficients were used evaluate the performance of the completed block of displays. If the participant's correlation coefficients was greater than 0.70, the display time for the subsequent blocks of trials was decreased. If at any point the correlation coefficient dropped below 0.50, time was added back to the display time. Participants oscillated between a maximum number of depth planes in the random paradigm between 6 and 8 total depth planes. Early learning consisted of more blocks of trials that consisted of a maximum of 6 depth planes, while later learning consisted of mostly 8 depth planes as maximum.

c. Results

From the approximately 10 originally naïve participants who have experience with the displays, only 3 were able to dedicate the time to decrease their display time to the desired 400 ms. The remaining 7 subjects, who participated for one or two hours, never accurately reported more than 3 depth planes unless given several seconds of

display time. There was no evidence that the learning demonstrated in the fixed depth plane experiment transferred to the randomized paradigm.

Since the goal of the experiment was to demonstrate how many depth planes may be seen simultaneously, these 3 dedicated participants showed that it is possible to accurately perceive up to 8 depth planes within 400 ms, thus presumably, all at the same time.

By utilizing *bits lost* (explained in the “Statistical Analysis” section), we were able to evaluate the visual learning that occurred over the course of more than 20 sessions, spanning approximately 3 months, in the 3 dedicated participants. Figure 4 depicts a comparison between blocks that has a maximum of 6 or 8 depth planes. The left portion of the graph shows blocks of trials where the *bits lost* values were substantially high (i.e. between -10 and -4 bits). As time and learning progressed, the *bits lost* per trial became substantially lower (i.e. between -1 and -0.3). It is important to note that while the *bits lost* values became smaller, the display time also gets smaller. The range of display time from 10 sec down to 0.4 sec.

Once learning stabilized, we were able to determine probabilities of each response given each stimulus. For fewer depth planes, reports were always accurate (Probability = 1, left-hand points in Figure 6). For more depth planes, report accuracy fell off with subjects sometimes reporting 1 less depth plane than presented (Probability < 1, right hand points in Figure 6). Accuracy for this, and other dedicated participants, was near perfect up to 6 depth planes and then began to fall off gradually. Other tests showed that these subjects could discriminate 6 depth planes from 7, and also 7 from 8, without error (in a Discrimination Experiment to be explained in a later section). Errors only occurred

when an unexpected number of depth planes were presented, and there were 6 or more of them.

V. Late Learning and Opposing-Contrast Displays: Experiment 3

a. Problem

Harris & Parker (1995) contended that opposing contrast displays will not produce the perception of depth. Based on the evident learning that occurred with standard RDS displays with white dots and a grey background, we postulated that a similar learning phenomenon may facilitate opposing contrast displays. Contrasting displays are defined as black dots in right visual field, white dots in left which stereoptically correspond. According to many studies (Harris & Parker, 1995; Smallman & McKee, 1995), the visual system should be unable to parse depth information since the conflicting signal will cause instability in the visual input; hence, the visual system cannot make sufficient matches in space. The reported sensation that optical focus cannot be established seemed to end scientific inquiry on the subject. No study has, to our knowledge, presented opposite contrast displays over a span of time to see if the perception of depth may be learned. During the Late Learning stages of the randomized paradigm, we began examining the opposing contrasts problem. We were also interested in seeing if learning from the standard randomization paradigm would transfer to opposing contrast displays.

b. Procedure

The MATLAB program was altered to produce corresponding points that would not match in optical contrasts (i.e. white dots in the left visual field that correspond with 50% black and 50% white dots in the right visual field; both fields being grey). All other aspects of the apparatus and program remain constant, including the priming display and

the use of the Brewster Method. Two dedicated participants, who already showed proficiency with the 6 and 8 depth plane randomized paradigm ($R^2 > 0.8$, *bits lost* < -2.0) when the dots were of the same contrast, were introduced to the new program. Due to the initial discomfort with the displays, blocks of trials were mixed with the standard randomization paradigm (6 or 8 depth plane maximum). We used the same *bits lost* analysis as described in the Statistical Analysis section to examine learning over time. The opposing contrast displays had a maximum of 6 depth planes when presented randomly.

c. Results

Contrary to the findings by Harris & Parker (1995), the human visual system is capable of acquiring depth information (up to 6 depth planes) within stimuli with same and opposite contrasts. As depicted in Figure 5, the learning curve of the standard randomization paradigm was similar in shape to the learning curve of the opposing contrast displays, despite the trials being months apart for the different experimental design. The left side of the graph shows the similar level of errors in the early stages of learning, but gradually the time interval and the errors made diminished over time. The right side of the graph (Figure 5) shows performance of both opposing contrast and same contrasts with a 400 ms display time and minimal errors made.

Similar to the lack of transference of skill from the fixed displays to the randomization paradigm, it appeared that learning for the distinctly different visual tasks may all begin at the same point; however, the data which indicate this are only complete for one participant, and so this conclusion needs verification by others.

VI. Discrimination: Experiment 4

a. Problem

Since accuracy on the 400 ms displays seemed to drop slightly after 6 depth planes, we wanted to investigate whether participants could in fact discriminate between two randomly displayed consecutive RDS stimuli (i.e. 6 depth planes from 7; 7 depth planes from 8; etc). The displays were programmed to present the participant with either n depth planes or $n+1$ in a random fashion, but like our standard randomized paradigm, each was presented an equal number of times per trial block. We presumed that if the dedicated participants could not discriminate consecutive displays when they knew that only those displays would be presented, they would not be able to report the number of depth planes in the randomization experiments.

b. Procedure

We altered the standard paradigm in order to give Late Learning participants a difficult discrimination task. We first employed the standard 3 pixel disparity used in all previous randomized and fixed displays. Participants had no difficulty discriminating 6 from 7 and 7 from 8 depth planes, when they knew that each was equally likely. This result is evidence that the drop-off in accuracy in the 2-6 or 2-8 experiments was not due to an inability to discriminate.

We then reduced the disparity from 3 pixels down to 1 pixel. Again, the displays were programmed to present the participant with either n depth planes or $n+1$ in a random fashion, but like our standard randomized paradigm, each would be presented an equal number of times per trial block. Performance was evaluated via the total *bits lost* per block of trials.

c. Results

With a high level of accuracy, participants could distinguish the difference between two depth planes (n , $n+1$) even with the more difficult disparity increments (1

pixel). However, similar to the randomized displays of up to 8 depth planes, all discrimination up to and including 6 depth planes was perfect or near perfect. One dedicated participant could distinguish 9 depth planes from 10 with a fair degree of accuracy (80%; chance being 50%). The fact that accuracy with the harder 1-pixel disparities began to drop off with 6 versus 7 depth planes indicates an underlying physiological limitation which could impact the data obtained with the easier 3-pixel disparities and 2-8 randomized depth planes. However, it is not known how disparity trades off against the information load (judging one of two possible depths versus one of 6 or 8), so this conclusion is speculative.

VII. Discussion

Supplementing previous research (Wilcox et al, 2008), the findings demonstrated that the human visual system, under optimal conditions and after learning, is capable of perceiving 6 or more depth planes simultaneously. Some participants may perceive up to 8 depth planes with a display duration of 400 ms, brief enough to disallow any vergence eye movements, and demonstrate that multiple depth planes may be perceived in unison. The number of trials at different stimulus time intervals varied per participant, which is expected. The phenomenon of visual learning that appears to facilitate the marked improvement in the task is difficult to account for, but the data strongly suggests that learning does occur. Approximately 10 participants have experience with the study. Those who participated for just one or two hours could not complete the task, scoring no better than chance. The dedicated participants who, over the course of weeks, practiced and gradually decreased the display time of the RDS show substantial improvement.

In the case of participant DL, the comparison between early trials and practiced trials with the RDS stimuli demonstrates the statistically significant change in

performance. In addition to improved accuracy, DL was about to decrease display time by a factor of more than 6, beginning with a 3 second display and eventually graduating to a 400 ms display. The theoretical lines used in the graphs are forced through (0,0). The x axis of the graph represents the number of depth planes presented by the randomized RDS stimuli, while the y axis represents the participant's response. The improvement from practice session (random performance) to stereoptic expert ($R^2 = .95$) demonstrates the importance of practice. The data suggests that with repeated exposure to RDS stimuli over months, a novice subject may be able to differentiate up to 6 to 8 depth planes in a randomized series with a high level of accuracy. Once the participant became a 'stereo expert,' he reported that he did not need to count the number of depth planes, it was more of an 'instinctual feeling' that resulted in correct responding. This was also the case with the other expert, MT.

Contrary to the findings by Harris & Parker (1995), the human visual system is capable of acquiring depth information (6 to 8 depth planes) within stimuli with same and opposite contrasts. Gepshtein and Cooperman (1998) utilized RDS displays in the shape of a cylinder. Participants were asked to decide between two orientations of the shape. They found that performance deteriorated when the density of the dots in the front of the display were increased. Performance improved when the dot density of the back were increased. With opposite luminance polarities (i.e. black vs. white), the researchers theorized that neurons in the human visual cortex that detect different disparities inhibit one another (Gepshtein & Cooperman, 1998). They were unable to discern why the opposite luminescence polarities cause such inhibition. Smallman and McKee (1995) discerned that a possible explanation for the contrast threshold within stereopsis is a biological safeguard against false matches and attendance to ambiguity. We found that

the brain's ability to organize perceptual items (even completely unrelated items) may trump that hardwired safeguard. The researchers went on to say that the result of the unmatched points resulted in the sensation of a 'ghostly feature' near the fixation point. Participants in our experiment reported a similar phenomenon, but as experience grew the perception of the 'ghostly figure' was not as prevalent. In both studies, the participants did not undergo a rigorous learning phase to acquire expertise of RDS displays. They could not clear the hurdle of the initial response of the visual system. Since we presented the opposing contrasts stimuli after participants were successfully parsing standard white on white RDS displays, future research should examine the trajectory of learning for the opposing contrast displays without prior RDS experience.

We have found that the learning necessary to successfully distinguish these novel depth displays in adults is slow and progressive, resembling the learning of depth cues in infancy. Even discerning depth information from 'unmatchable' stereoptic displays may be facilitated by learning. The overall plasticity of the visual system should not be surprising; slight alterations and modification may occur consistently through our lives when we encounter new visual experiences. With the future of entertainment and media headed into the realms of 3 dimensions, such finding may guide development and optimalization of these technologies.

Figures

Figure 1. “Binocular disparity.” With the slightly different inputs from each eye, the visual system utilizes corresponding points to produce the percept of depth.

Figure 2. “Dramatization of RDS display with 4 transparent depth planes.” The adaptation shows the pile of transparent depth planes that the computer manipulates to produce the sensation of depth. The arrow represents the direction of viewing.

Figure 3. “The apparatus.” Picture ‘a’ shows a participant seating in the experiment. In the image, one can see the specula, computer screen and a sample display. The Brewster Method, using a specula and a prism lens to direct the line of site, was used to produce the illusion of depth. Picture ‘b’ is a photograph of the computer screen during a trial. Note the two distinct images that are directed to each visual field.

Figure 4. “Learning Curve in Bits Lost: 2-6 depth planes vs. 2-8 depth planes, random, WW, 400 ms (3 step moving average: S=D. Lynch).” The graph shows the learning curve, presented as *bits lost*, for randomized 6 and 8 maximum depth plane trials (both visual fields receive white dot stimuli). Each major point on the graph represents a single trial of randomized displays.

Figure 5. “Learning Curve in Bits Lost: 2-6 depth planes, WW vs. BW, random, 400 ms (3-step moving averages: S= D. Lynch).” The graph compares the learning curves between white-white and white-black stimuli. According to Harris & Parker (1995), it is not possible to perceive depth with opposing contrasts. We have demonstrated that it is

indeed possible to, with a high level of accuracy, differentiate up to 6 depth planes with white-black contrast. Note the similarities between the learning curves of white-white and black-white.

Figure 6. “Probabilities of Responses: 2-8 Depth Planes, WW, Random, 400 ms (S=M. Tran).” The graph shows the probability that one participant responds with each depth plane value. The figure depicts results during and after learning has stabilized. This data shows the number of depth planes perceived simultaneously since the time interval does not allow for vergence eye movements.

References

- Berkeley, G. (1709) *An Essay Towards a New Theory of Vision*. Reprinted: London, Everyman's Library, 1989.
- Blakemore, Colin. Hague, Brian. (1972) "Evidence for disparity detecting neurons in the human visual system." *Journal of Physiology*. Vol. 225: 437-455.
- Falk, David S. Williams, Robert. (1980) "Dynamic visual noise and the stereophenomenon: Interocular time delays, depth, and coherent velocities" *Perception & Psychophysics*. Vol.. 28.1: 19-27.
- Gepshtein S.; Cooperman A. (1998) Stereoscopic transparency: a test for binocular vision's disambiguating power. *Vision Research*: 38, 19, October: 2913-2932(20).
- Harris, J. Parker, A. (1995) Independent mechanisms for bright and dark information in binocular stereopsis. *Nature*. 374: 808-811.
- Julesz, B. (1973) "Cyclopean perception and neurophysiology." *Investigative Ophthalmology*. Vol. 11:6, June.
- McKee, S.P. Verghese, P. (2002) "Stereo transparency and the disparity gradient limit." *Vision Research*. 42: 1963-1977.
- Ogle, K.N. (1950). *Researches in binocular vision*. Philadelphia: W.B. Saunders Co.

- Satgunam, P., Gowrisankaran, S., & Fogt, N. (2009). The influence of vergence adaptation on open-loop vergence dynamics. *Vision Research*, 49, 1795-1804.
- Semlow, J. L. & Wetzell, P. (1979). Dynamic contribution of the components of binocular vergence. *Journal of the Optical Soc. of Am.*, 69, 639-645.
- Smallman H.S., Mckee S.P. (1995) A Contrast Ratio Constraint on Stereo Matching. *Biological Sciences*; 260, 1359, June, 265-271.
- Weinshall, D. (1989) "Perception of Multiple Transparent Planes in Stereo Vision." *Nature*. Vol. 341. October.
- Wilcox, L. Tsirlin, I. Robert, A. (2008) "Stereoscopic transparency: Constraints on the perception of multiple surfaces." *Journal of Vision*. 8.5: 1-10.
- Zabulis, X. Backus, B. (2004) "Starry Night: A Texture Devoid of Depth Cues." *Journal of Opt. Soc. Am*: 21, No. 11, November.

Figure 1

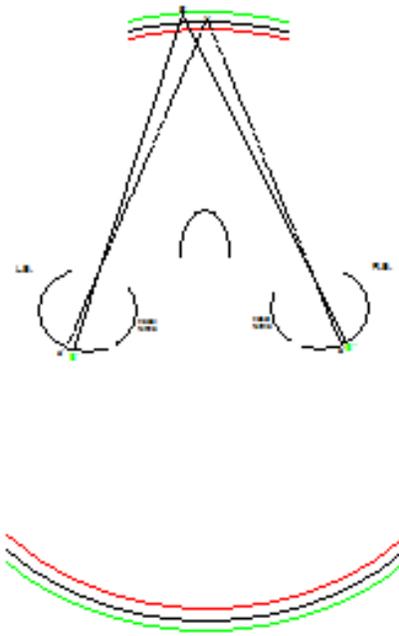


Figure 2

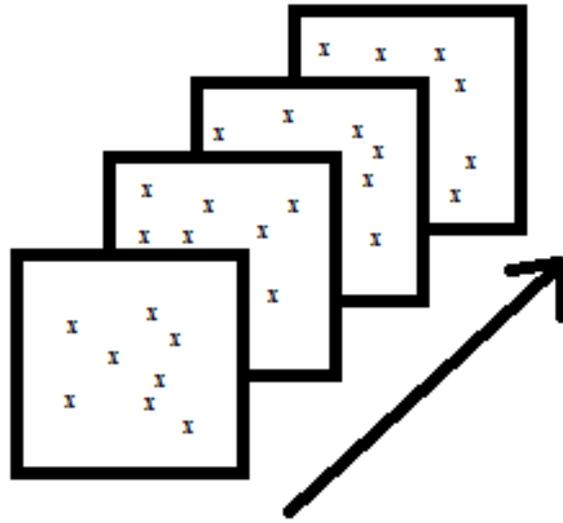
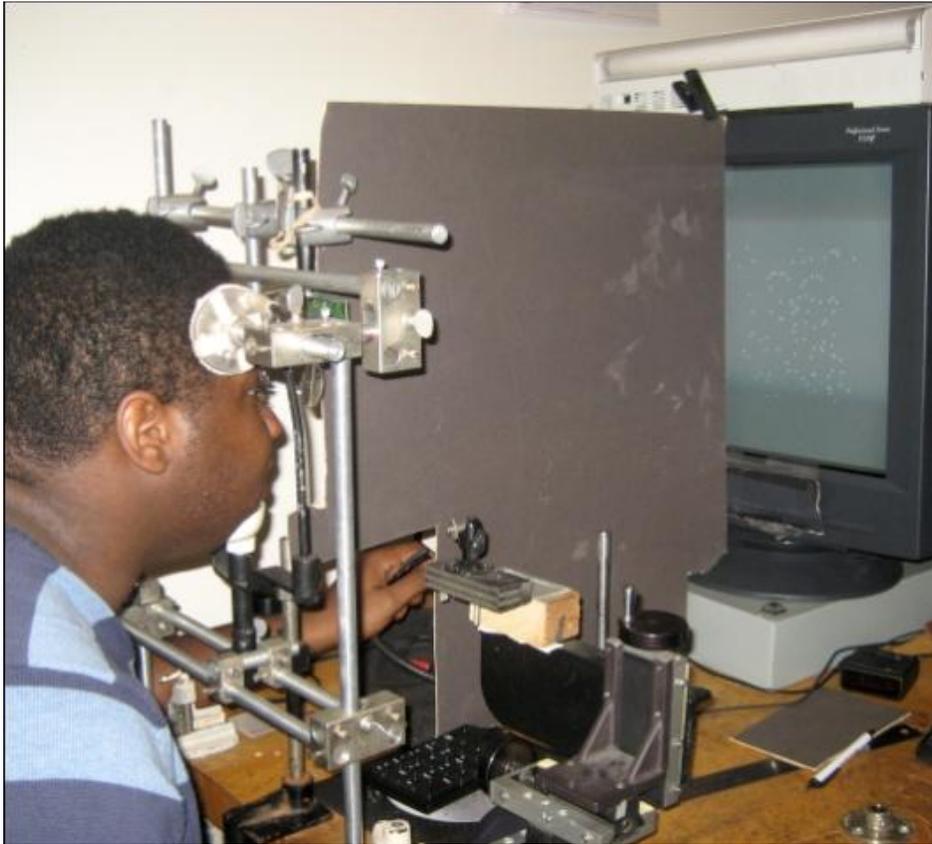


Figure 3

a.



b.



Figure 4

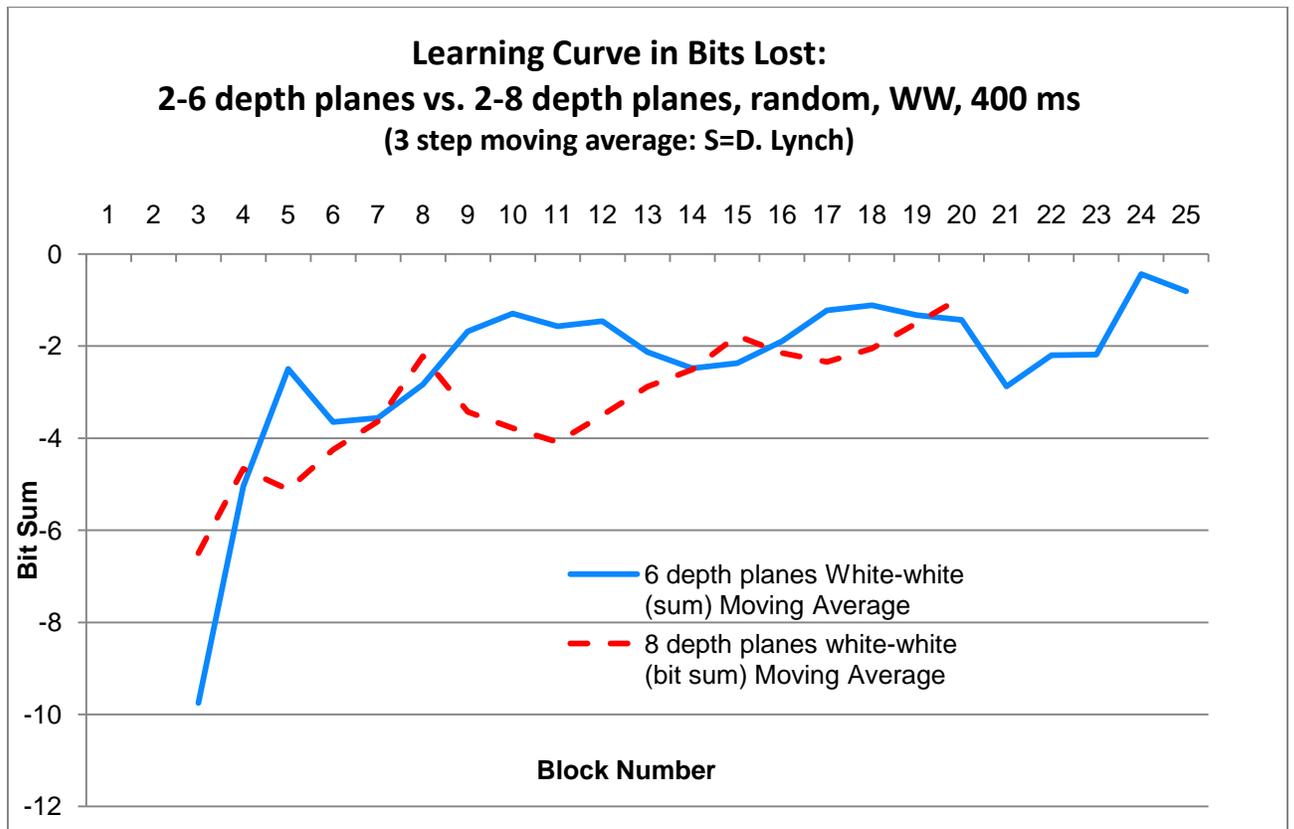


Figure 5

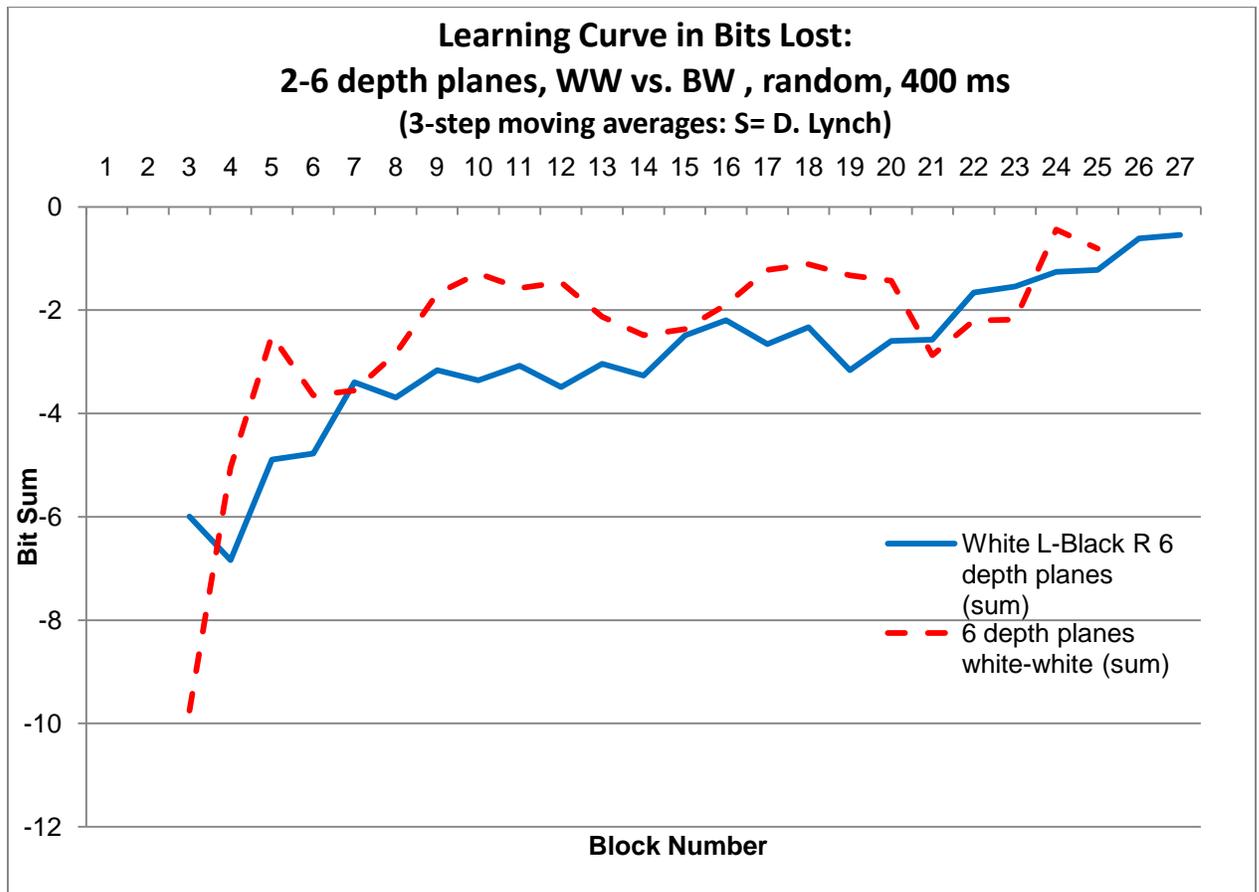


Figure 6

