The Role of Contrast in the Perceived Depth of Monocular Imagery

Allan G. Rempel* — The University of British Columbia
Wolfgang Heidrich — The University of British Columbia
Rafał Mantiuk — Bangor University

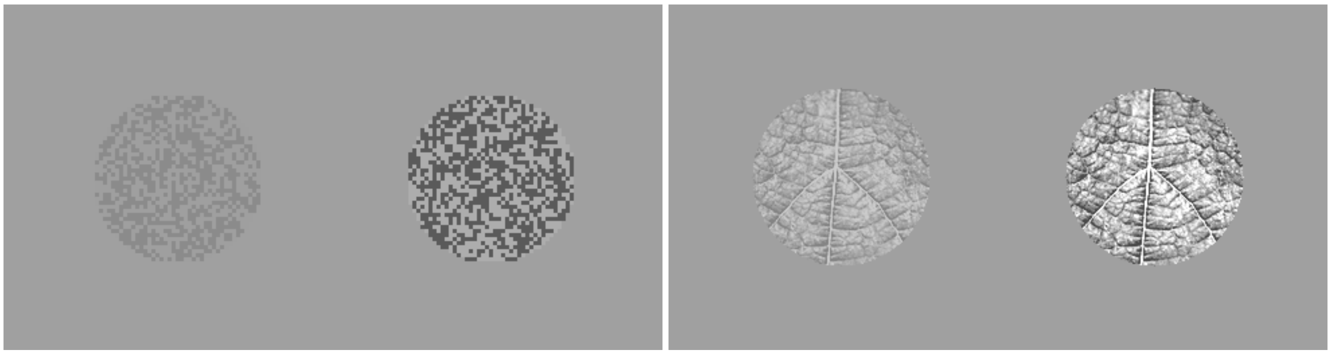

Figure 1: *The perceived depth of an object is related to the contrast both between the object and the background, and within the object itself.*

Abstract

Since high dynamic range (HDR) displays have been shown at conferences, they have been confused with 3D displays by some observers. In this paper, we explore this perceptual connection by conducting a series of experiments to examine the effect that contrast has on depth perception. In particular, we consider the contrast both of large-scale features and of small-scale features, both independently and in concert. We found that in each of three experiments, subjects perceived increases in contrast to correspond with increases in perceived depth. Our findings indicate that we can simulate sensations of depth by manipulating contrast, particularly that of highlights within images, and that modern high-contrast displays can simulate greater sensations of depth.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.3.6 [COMPUTER GRAPHICS]: Methodology and Techniques—Interaction techniques.

1 Introduction

Since ancient times, scientists have considered the cues that provide us with sensations of depth as the light from three-dimensional scenes is projected onto our curved two-dimensional retinae. These phenomena have been well known already for many decades [Boring 1942, ch. 8].

Optical illusions provided some early opportunities to study the different types of cues and observe when they produce accurate 3D percepts and when they do not. For instance, the *Moon Illusion*, in which the size of the moon at the horizon appears larger than when higher up in the sky, was known to the ancient Greeks. von Helmholtz [1924-25, v. III, p. 360-362] discusses the history of explanations of this phenomenon from Ptolemy (150 A.D.) through to his time, with the conclusion being that the intermediate objects along the horizon between the viewer and the moon lead the viewer to perceive the moon to be farther away at the horizon than it is higher up, and since it occupies the same angular size, it is perceived to be larger.

New technologies can also expose characteristics of depth perception. For instance, one can observe a difference in the strength of 3D percepts between television programs viewed on smaller and larger displays, or between standard-definition and high-definition displays. Sports telecasts such as hockey and football games show the difference particularly well. This suggests that depth perception may be related to display characteristics such as resolution, brightness, contrast, or size. These characteristics can reduce the gap between the retinal image produced by the display and that which would have been produced by viewing the actual scene.

Some depth cues such as *parallax* and *oculomotor* effects can be perceived only by viewing the actual 3D scene. Parallax refers to having two different views of a scene, and can come from motion (of the object or the observer) or from the *stereopsis* feature of human binocular vision, which is the ability to fuse two slightly different retinal images into a single 3D image.

In addition, there are many other, generally monocular, depth cues that may be captured in a painting or photograph or other 2D image of a 3D scene, such as perspective, relative sizes of objects, familiarity with sizes of objects, occlusion, and *aerial perspective* which includes contrast, color saturation, and haze. Contrast and brightness are of particular interest to us since they can be manipulated through a much greater range on high dynamic range displays than is possible on conventional displays.

High dynamic range (HDR) displays such as those demonstrated by Seetzen and colleagues [2004] have a dramatically increased dynamic range over conventionally available *low dynamic range* (LDR) display technologies, and come much closer to meeting the abilities of the human visual system (HVS). In demonstrations of these displays, many observers have remarked that they appear to be able to represent 3D to a much greater degree than conventional displays [Seetzen 2007]. In this work, we examine some facets of 3D (depth) perception and conduct psychophysical experiments which examine the nature of this depth perception and why the effect is so much more prominent on HDR displays.

2 Related Work

One of the first descriptions of contrast as a depth cue dates back to Leonardo da Vinci, who observed the phenomenon of aerial perspective, in which atmospheric haze reduces the contrast of distant objects and alters the color to be more blue [Richter et al. 1939, p. 210-212, 234-241].

*email: agr@cs.ubc.ca

More recently, many experiments have been conducted to determine the effects of brightness and aerial perspective on depth perception.

Brightness. Miles [1953] observed the interaction between brightness and depth perception by noting that binocular viewing of two images with unequal brightnesses altered the perception of the depth of objects. Gilchrist [1980] conducted an experiment in which patches of paper at varying depths and with varying luminance were seen monocularly through a pinhole, and found that the perceived lightness (from a Munsell chart) of a constant-luminance patch differed dramatically depending on how far away the patch was perceived to be. Gilchrist also suggested that the causality between perceived lightness and perceived depth may be bidirectional. Schirillo et al. [1990] built on Gilchrist's experiments, separating the concepts of brightness and lightness, and using binocular stereoscopic viewing of a computer monitor in place of the physical setup Gilchrist had used.

Contrast. O'Shea et al. [1994] investigated aerial perspective by conducting an experiment which showed on a computer monitor that objects with higher-contrast edges are deemed to be nearer to the viewer than objects with lower-contrast edges. Rohaly and Wilson [1999] conducted a set of experiments in which they found a power-law relationship between contrast and perceived depth using stimuli consisting of vertical bars with Gaussian luminance profiles. Ichihara et al. [2007] extended this work by differentiating area contrast (that of large-scale features) from texture contrast (that of small-scale features) and observing that both have an effect on perceived depth. And Fattal [2008] used a model of aerial perspective to construct depth maps of elements within scene images in order to reduce the haziness within those images.

Color. Triesman [1962] observed that both contrast and color could affect depth cues in the viewing of stereoscopic scenes. More recently, Troscianko et al. [1991] observed that while a color gradient between red and green did not significantly affect depth, a color saturation gradient between red and grey did.

Integration of Depth Cues. The various different depth cues each have their effects, but the relative strengths of those effects can vary, and that variance is often based on image content. Wijntjes and Pont [2010] observed that binocular stereo could improve the depth perception in images, but it is not a foregone conclusion that stereo is always better as not all images in their study benefited from a stereo representation. Held et al. [2010] manipulated focus and blur in a manner similar to tilt-shift photography to dramatically change the perception of the scale of a scene. And Cipiloglu et al. [2010] developed a framework based on fuzzy logic for enhancing depth perception of imagery using many different depth cues.

3 Experiments

We conducted three experiments to analyze the relationship between contrast and perceived depth of objects in an image. Our first experiment was designed to reproduce the results obtained by Ichihara [2007] on random dot patterns (Section 3.1), and extend them to natural textures (Section 3.2). Our second experiment (Section 3.3) was designed to analyze the impact of a higher dynamic range on depth perception. Finally, our third experiment (Section 3.5) analyzes the impact of tone curves on the depth perception in natural scenes.

3.1 Experiment 1a: LDR Texture Contrast and Area Contrast

The experiments of [Ichihara et al. 2007] are well designed to show the effect of brightness and contrast on depth perception in the context of conventional LDR displays. These are the depth cues upon which HDR technology can have the most significant effect, so they will be the focus of this study. Our first experiment mimics closely their Experiment 2, with the exception that ours was conducted on an HDR display simulating LDR by using a uniform backlight, while theirs was conducted on a conventional LDR display.

Subjects. Ten subjects (8 male, 2 female, aged 22–42) participated in Experiments 1 and 2. Each of them had normal (20/20) or corrected-to-normal vision, which was confirmed through the administration of a Snellen visual acuity test. Each session lasted approximately one hour, after which subjects filled out a questionnaire about their experience and perceptions.

Stimuli. The stimuli consisted of random-dot disks, shown horizontally side-by-side, on a uniform background. Some examples of random-dot disks are shown on the left side of Figure 1. The distance between the centers of the disks was $5.7°$ of visual angle, and each disk had a diameter of $3.2°$ (132 pixels). Each disk was composed of an equal number of light and dark dots, arranged randomly throughout the disk. Each dot was a 3×3 pixel square. The stimuli were presented on a Dolby DR-37P HDR display at a viewing distance of 100 cm, at eye level in the center of the subject's field of view.

Each pair of disks included a reference disk and a test disk. The reference disk had a constant texture contrast of 0.5 and a constant average luminance of 30 $\mathrm{cd/m^2}$. For the test disk, texture contrast was varied from 0.1 to 0.9 in steps of 0.1, and the average luminance was varied from 10 to 50 $\mathrm{cd/m^2}$in steps of 10 $\mathrm{cd/m^2}$. For this experiment, the Michelson contrast metric was used. Texture contrast was defined as the contrast between light and dark dots within a disk ($(L_{light} - L_{dark})/(L_{light} + L_{dark})$), while area contrast was defined as the contrast between the average luminance of a disk and the uniform background luminance ($(L_{disk} - L_{background})/(L_{disk} + L_{background})$). Background luminance levels of 20 and 40 $\mathrm{cd/m^2}$were used.

Experimental procedure. Subjects were shown a series of pairs of random-dot disks, and asked to indicate on the keyboard for each pair which disk appeared closer (or larger) and by how much. A 7-point scale was used to encode whether subjects found the left or right disk to be closer by a low, medium, or high amount, or whether they were perceived to be at about the same depth.

A practice run of 10 pairs of disks covering a representative range of contrasts was first conducted to familiarize subjects with the mechanics of the experiment. This was then followed by 90 trials: 2 background luminance levels $\times$ 5 average disk luminance levels $\times$ 9 texture contrast levels. The full set of 90 conditions was randomized, and the positions of the reference and test disks were exchanged at random. Subjects went through the experiment at their own pace.

Results. The results are substantially similar to those of Ichihara et al. Increases in texture contrast corresponded with increases in the perceived closeness of a disk, as shown in the left plot of Figure 2. This corresponds with Figure 4 in [Ichihara et al. 2007, p. 692]. Where Ichihara's results are broken down into two plots (for background levels of 20 and 40 $\mathrm{cd/m^2}$), we have omitted some

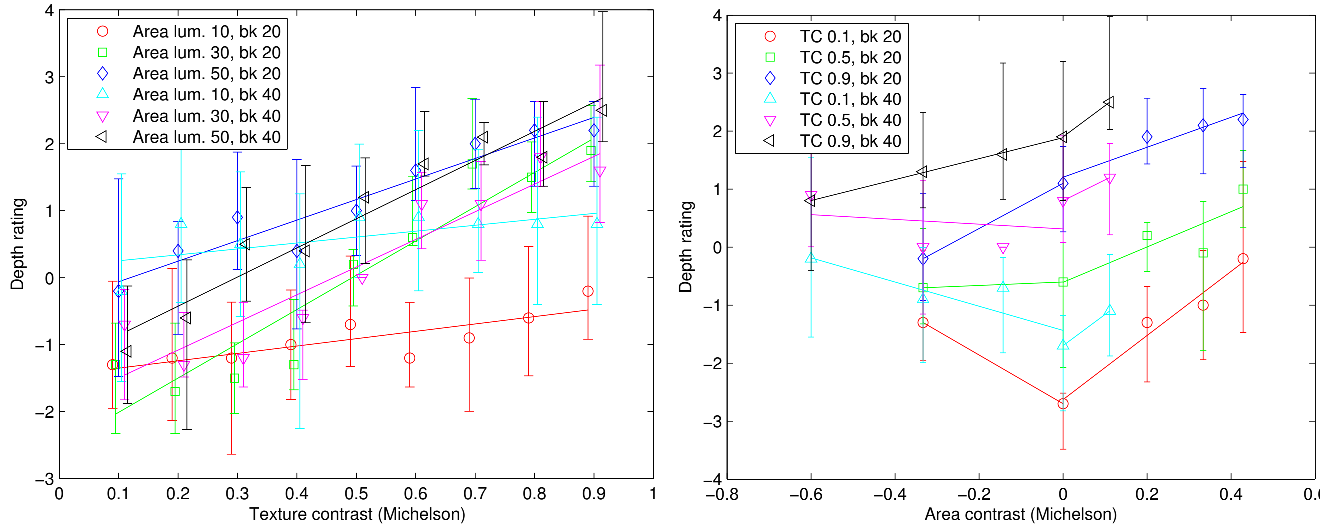

Figure 2: *Random-dot disk images: Mean depth judgements as a function of texture contrast (left) at different levels of average disk luminance (10-50 cd/m^2), and area contrast (right) at different levels of texture contrast (0.1-0.9), at two levels of background luminance (20,40 cd/m^2). The solid lines show best fit using least-squares approximation, while the error bars show the standard deviations at each point. The points are offset slightly to improve clarity.*

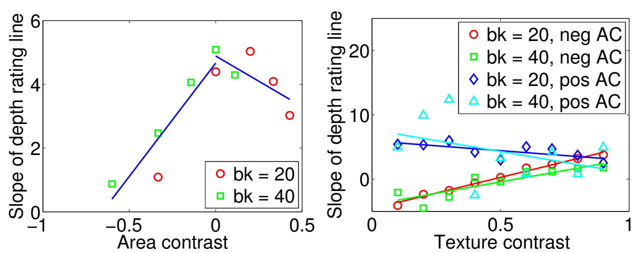

Figure 3: *Random-dot disk images: The relationship between area contrast and the degree to which texture contrast affects depth perception (left), and the relationship between texture contrast and the degree to which area contrast affects depth perception (right).*

intermediate data and combined the results into a single plot for clarity and brevity. The full results are available at our project web site[1]. Further, the effect of texture contrast was significantly more pronounced at low levels of area contrast than at high levels, as shown in the left plot of Figure 3. This corresponds with Figure 5 in [Ichihara et al. 2007, p. 692].

We also used the experimental data to analyze two questions not addressed by Ichihara et al.: how does depth perception vary with area contrast at constant texture contrast levels, and how does varying the texture contrast affect those curves? The right plot of Figure 2 shows that perceived closeness generally increases as the area contrast of a disk increases, with texture contrast being held constant, when the area contrast is positive (bright area on a dark background). When the area contrast is negative (dark area on a brighter background), perceived closeness increases as the area contrast increases at low levels of texture contrast, but decreases as area contrast increases at high levels of texture contrast. The effects of increases in texture contrast, mitigating the effect of positive area contrast and reversing the effect of negative area contrast, are shown in the right plot of Figure 3.

[1]http://www.cs.ubc.ca/labs/imager/tr/2011/Rempel_Depth_Perception

We found the correlation between depth rating and texture contrast to be statistically significant in all cases except those with the lowest area luminance (10 cd/m^2) as indicated by the $F_{8,81}$ values in the upper ("dots") row of Table 1. The same is true for the "leaf" row which shows the results for Experiment 1b (Section 3.2). We also found statistical significance for the correlation between depth rating and area contrast, particularly for those line segments in the right plot of Figure 2 that are defined by four points (as opposed to those defined by two); the $F_{3,36}$ values are given in the bottom row of Table 1. Ichihara et al. found statistical significance in the difference between the slopes of the lines relating depth rating and texture contrast (the left plot of Figure 2) and so did we ($F_{4,45} = 21.10$ (dots) and $F_{4,45} = 10.48$ (leaf) at 20 cd/m^2, $F_{4,45} = 14.20$ (dots) and $F_{4,45} = 9.44$ (leaf) at 40 cd/m^2, $p < 0.01$).

Back lum.	20			40		
Area lum.	10	30	50	10	30	50
dots	–	55.39	9.05	–	31.55	20.25
leaf	–	62.90	26.70	3.27	70.31	38.43
Tex con.	0.1	0.5	0.9	0.1	0.5	0.9
dots	10.93	5.23	5.50	$3.28^{\dagger}$	$3.06^{\dagger}$	–
leaf	$3.70^{\dagger}$	28.95	–	7.45	7.80	–

Table 1: *F values from ANOVA performed on depth rating vs. texture contrast at varying area luminance (above) and depth rating vs. area contrast at varying texture contrast (below), with $p < 0.01$, or $p < 0.05$ where denoted by (†). Some results, denoted by "–", did not reach statistical significance.*

3.2 Experiment 1b: Natural Textures

For the purposes of seeing how contrast affects natural scenes, it is useful to know whether these relationships extend to images of naturally occurring texture in addition to random-dot images. We tested this by running the same experiment as above, but using a leaf texture adapted from a close-up photograph instead of the random-dot images. For convenience, we integrated the two textures (dot and leaf) into the same experimental run, doubling the total number

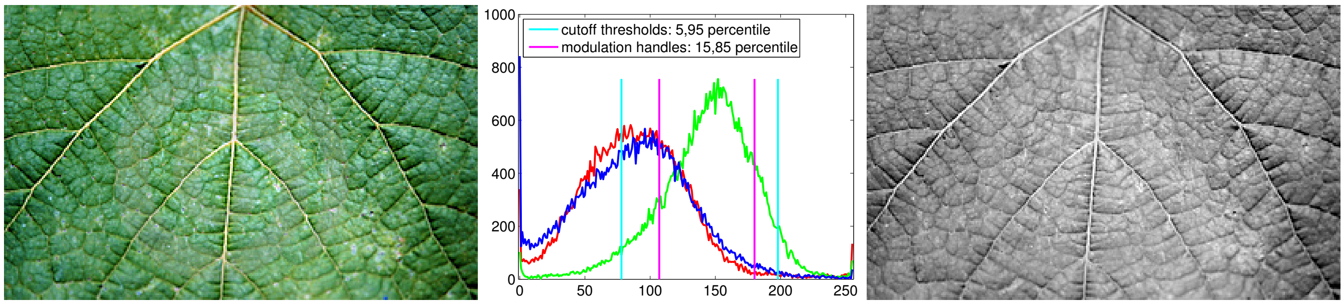

Figure 4: *Left: the original leaf color photograph image; Middle: the histograms for R, G, and B, with thresholds based on G; Right: the desaturated image where $R = G$ and $B = G$.*

of trials from 90 to 180, where each trial could (at random) be either a leaf or dot texture. Therefore, the same subjects participated in both parts of the experiment.

Our experiments in this work relate to intensity and contrast, but not color, so it was necessary to desaturate the color from the leaf image. There are a variety of standard techniques for this, such as the luma computations of ITU-R Recommendations BT. 601 ($Y = 0.299R + 0.587G + 0.114B$) or BT. 709 ($Y = 0.2126R + 0.7152G + 0.0722B$), but since the leaf image was predominantly green, we used the simpler technique of setting the red and blue channels to be equal to the green channel (i.e. $Y = 0R+1G+0B$). The differences between these techniques would normally manifest themselves in brightness and/or contrast differences, but since our experiment controls both of those characteristics, there should be no difference in our experiment between those techniques. To confirm that, we tested all three techniques under a broad range of experimental conditions and observed no visible differences between the three techniques.

The histogram of the image shows a roughly Gaussian shape as one would expect, which is significantly different from the binary distribution of the random-dot pattern. We adjusted the distribution of the leaf image to make the texture contrast and overall area luminance consistent with that of the random-dot distribution through a two-step process. First, we clamped pixel values that exceeded low and high thresholds which we defined by the 5 and 95 percentile levels of the image histogram (i.e. of the green channel). This prevented outlier pixels (such as the unnaturally large number of clamped pixels at 0 and 255) from unduly increasing the contrast of small high-frequency details to a level out of proportion with the random-dot patterns at the corresponding texture and area contrast levels. Second, we established handles at the 15 and 85 percentile levels of the image histogram, which we moved (along with the image pixel values) to match the low and high levels of the binary distribution in the dot images in the various texture and area contrast configurations.

Multiple visual comparisons (over a broad range of contrast settings) between the random-dot and leaf images at the same contrast levels showed them to have similar levels of overall brightness and contrast. In addition, the measured luminance levels of the leaf images at different settings were commensurate with the measured luminance levels of the random-dot images. We also made ramp images consisting of all the pixels in a leaf image, sorted by value. That enabled us to measure the luminance of the dark and light areas of that leaf image, which were also commensurate with the measured luminance of dark and light areas of the corresponding random-dot images.

Figure 4 shows the leaf texture photograph [Andreas 2009], the histograms of the red, green, and blue channels, and the monochromatic version we created to use in our tests.

Results. The relationships between area contrast, texture contrast, and depth were substantially the same as those observed in the random-dot experiment, shown in Figures 2 and 3. The specific results are omitted here for brevity, but are available at our project web site. The results were statistically significant, as discussed in Section 3.1.

3.3 Experiment 2: HDR Texture Contrast and Area Contrast

Our next experiment was designed to analyze how the results of the previous section change when the range of contrasts is stretched out of the LDR domain and into the HDR domain.

The design of currently available HDR displays is based on the principle of dual modulation, which uses both high-resolution and low-resolution modulators [Seetzen et al. 2003]. A consequence of this design is that high dynamic range is only available with images that contain blocks of thousands or more contiguous pixels of similar brightness. In the context of the Ichihara experiments, it is therefore impossible to achieve texture contrasts beyond those available on LDR displays. However, high area contrasts can easily be achieved, and this is a significant component of the increased depth perception reported on HDR displays. However, as we will see in Section 3.5, it only accounts for part of the overall effect.

Subjects. The same subjects participating in Experiment 1 also participated in Experiment 2, which was conducted in the same session.

Stimuli. The stimuli used for this experiment were similar to those of the previous experiment. Again, random-dot disks and leaf texture images of the same size were used, but this time the distance between the disks was increased to 8.5° to allow the disks to be maximally coincident with a largely circular block of LEDs. The dual modulation scheme of the Dolby DR-37P HDR display that was used for this experiment employs an LCD panel and a backlight comprised of a hexagonal grid of LEDs. We placed the disks in order to fit as closely as possible to a block of 3 rows of LEDs with 2 LEDs in the top and bottom rows and 3 LEDs in the middle row. The grouping together of LEDs allowed for a greater range of area contrasts.

The range of luminances and contrasts was significantly greater than that used in Experiment 1. Background luminances remained at 20 and 40 $\mathrm{cd/m^2}$, but disk area luminances were extended to a range from 5 to 700 $\mathrm{cd/m^2}$using texture contrast levels of 0.2, 0.5,

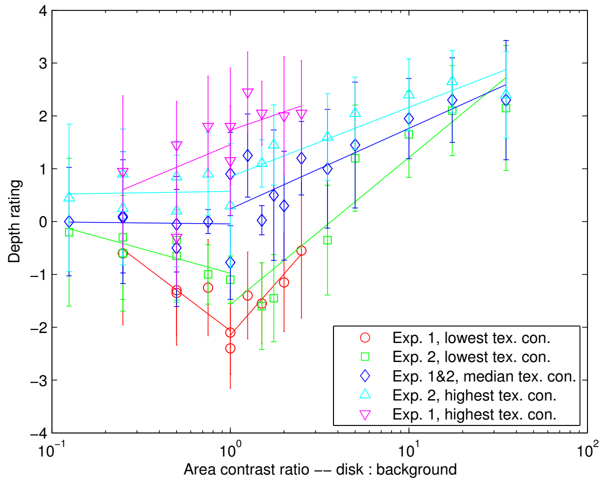

Figure 5: *The effect of varying area contrast ratios on the perceived depth of disks, at low, medium, and high levels of texture contrast. The error bars show the standard deviations around the means, while the solid lines are best-fit lines for their corresponding data points.*

and 0.8. The reference disk remained at a constant texture contrast of 0.5 and a constant average luminance of 30 cd/m^2.

Experimental procedure. The procedure was very similar to that of the previous experiment. There were a total of 54 trials: 3 background luminance levels $\times$ 3 average disk luminance levels $\times$ 3 texture contrast levels $\times$ 2 types of stimuli (random-dot disk and leaf image). Again, the full set of conditions was randomized, the positions of the reference and test disks were exchanged at random, and subjects completed the experiment at their own pace. The experiment was preceded by a practice run of 8 pairs of disks.

Results. The results of this experiment generally showed the same relationships between area contrast, texture contrast, and perceived depth as were observed in Experiment 1. Increases in texture contrast generally corresponded with increases in the perceived closeness of a disk, more so at low area contrast levels than at high (positive or negative) area contrast levels. Additionally, the dramatic increase in positive area contrast yielded a dramatic increase in the perceived closeness of a disk, particularly at lower texture contrast levels. The full results are omitted here for brevity, but are available at our project web site.

Figure 5 shows the effect of area contrast on perceived depth, and the increased effect of the high levels of area contrast achievable on HDR displays over the lower levels of area contrast achievable on LDR displays. In this plot, we use the contrast ratio metric (L_{light}/L_{dark}) which is better able to represent higher dynamic ranges than is the case with Michelson contrast. The plot includes the results using both the leaf textures and the random-dot textures, since the results from the two types of textures were very similar to each other. We found statistical significance for most of these relationships, as shown in Table 2. Additional results are available at our project web site.

	$CR \leq 1$	$CR \geq 1$
Exp. 1, lowest tex. con.	$F_{3,116} = 14.90$	$F_{4,115} = 13.10$
Exp. 2, lowest tex. con.	$F_{4,135} = 3.13^{\dagger}$	$F_{7,152} = 65.20$
Exp. 1&2, med. tex. con.	–	$F_{10,269} = 25.32$
Exp. 2, highest tex. con.	–	$F_{7,152} = 20.95$
Exp. 1, highest tex. con.	$F_{3,116} = 6.22$	$F_{4,115} = 3.79$

Table 2: *F values from ANOVA performed on depth rating vs. contrast ratio between the disk area luminance and the background luminance, with $p < 0.01$, or $p < 0.05$ where denoted by (†). These correspond to the plots in Figure 5. Some results, denoted by "–", did not reach statistical significance.*

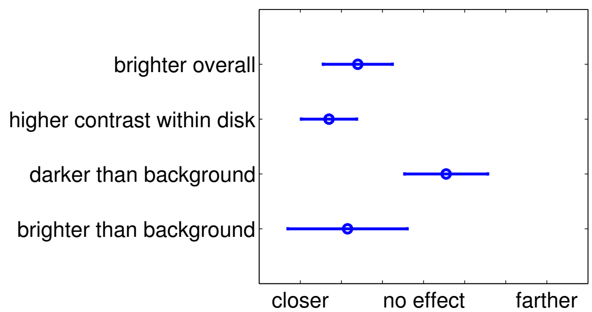

Figure 6: *The degree to which subjects perceived disks of different characteristics to be closer or farther.*

3.4 Questionnaire Results

In the questionnaire, subjects were asked whether they observed a relationship between contrast and the perceived depth of objects in the images, by marking on a 7-point Likert scale the degree to which they thought disks of different characteristics were closer or farther. Figure 6 shows the characteristics and the subjects' responses. Note that subjects indicated strongly that they perceived brighter disks and higher-contrast disks to be closer than others. They also indicated very weakly that disks that were darker than the background seemed farther. This qualitative observation corresponds well to the quantitative observation in the right plot of Figure 3. We surmise that the weakness of this indication is likely due to a possible ambiguity about the meaning of "darker"; it could refer to dark disks with a high contrast against the background which would seem closer, or disks that are darker than the (closer-seeming) bright disks and hence seem farther.

Subjects were also asked (on a 7-point Likert scale) whether they found it easier to make depth judgements with the random-dot disks or with the natural-image leaf-patterned disks, and indicated that it was much easier with the leaf-patterned disks.

3.5 Experiment 3: Contrast Modulation in Natural Scenes

The previous experiment has demonstrated the importance of area contrast for depth perception. In the next experiment, we analyzed the role of texture contrast for scenes with an already high area contrast. In particular, we observe how modulating the contrast of a texture patch through the application of a tone curve can affect the perception of depth within that scene.

Subjects. Eleven subjects (all male, aged 22–34) participated in the experiment. Two of them had been subjects in the first two

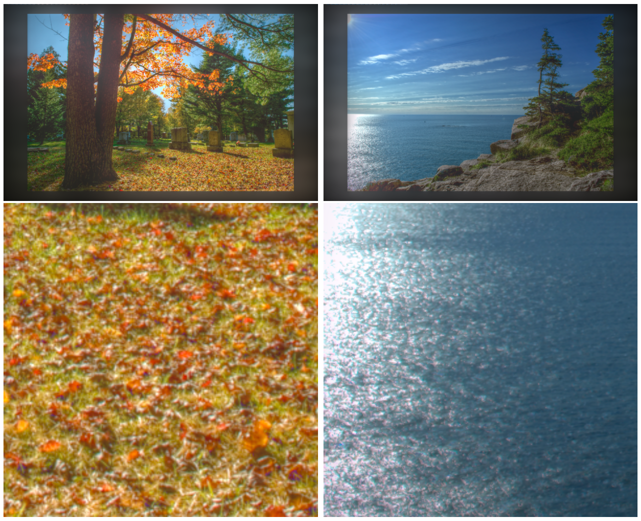

Figure 7: *Top: The full-sized HDR images. Bottom: The regions of the above images that were used in Experiment 3.*

experiments while the remaining nine had no prior experience with this research. Each of them had normal (20/20) or corrected-to-normal vision, which was confirmed through the administration of a Snellen visual acuity test.

Stimuli. The stimuli in this experiment were textured regions that had been cropped from two HDR images. The purpose of working with cropped regions rather than full images was to eliminate depth cues other than contrast, including large-scale shadows, occlusions, scale recognition from familiar large-scale features, and so forth. Figure 7 shows the larger images (above) and the cropped regions (below). The images were displayed on the same HDR display that was used in the previous experiments.

Each image was rendered at multiple different contrast levels in which the intensity of both the brighter and darker parts of the image were independently modulated. Within the texture patches, the mid-tones, defined as the intensities in the range of the median input level in the image, were kept relatively constant throughout these contrast manipulations. The patches where shown against a dark background, with the goal of keeping area contrast uniformly high across the experiment. Sigmoidal functions similar to the ones used by Mantiuk and Seidel [2008] were used to reduce the contrast in the dark and bright regions without causing discontinuities in the tone curve. Sigmoid or S-shaped curves have been commonly used in traditional photography to adjust the intensity distributions at both the high and low ends while avoiding intensity discontinuities [Reinhard et al. 2002]. The upper intensities were scaled to four different levels while the lower intensities were scaled to three different levels. Figure 8 shows logarithmic plots of the tone curves that were used for the two images (Leaves and Ocean) in this experiment.

Experimental procedure. Subjects were shown pairs of images in which the two were different only in the scaling levels of the upper and lower intensities. For each pair, the subject was asked to indicate on the keyboard which of the two showed a greater sense of depth within the scene. The screen was then cleared for 0.3 s before the next pair was displayed.

A practice run of 8 pairs of images selected at random from the full set was first conducted to familiarize subjects with the mechanics

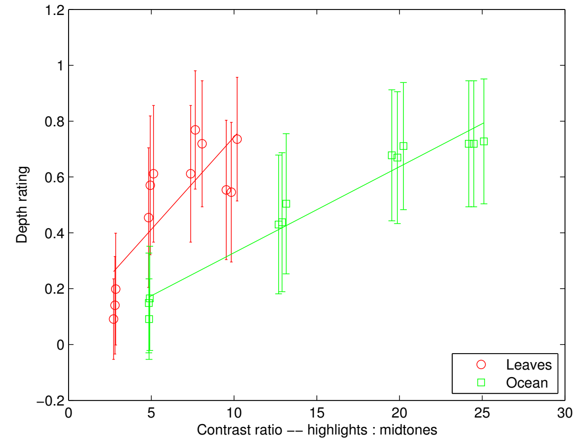

Figure 9: *Subjects' perception of the comparative depth of scenes with different levels of highlight contrast. The depth rating indicates the probability (based on our experimental results) that a viewer would identify that image as having greater depth than other images within the range of our study.*

of the experiment. This was then followed by 156 trials: all combinations of 12 contrast variations (4 upper contrast levels $\times$ 3 lower contrast levels) of an image taken 2 at a time ($C_2^{12} = 66$), plus 12 pairs where the same image was present on both sides, times 2 images. Subjects went through the experiment at their own pace, and there was a two-minute break at the half-way point of the experiment. The whole experiment took less than 25 minutes.

Results. Figure 9 shows the subjects' perception of the comparative depth of scenes with different levels of contrast between the highlights and the mid-tones. As in Experiment 2, we use the contrast ratio (L_{light}/L_{dark}) for our contrast calculations. Each point in this plot represents 1 of the 12 contrast variations for an image, where the horizontal axis represents the ratio between the bright point and the mid-tone for that variation. The vertical axis represents the mean depth rating score, relative to the other versions of the image; the error bars indicate the standard deviations. The vertical range is normalized to (0,1), where 0 indicates that all other variations across all subjects were seen to have more depth than the current variation, while 1 indicates that all other variations were seen to have less depth, and 0.5 indicates that an equal split between the number of images that were perceived as containing more vs. less depth. As indicated in the plot, subjects indicated a greater sense of depth as the contrast of the highlights was increased, for both images. These results were statistically significant for both the Leaves ($F_{11,714} = 14.57$) and Ocean ($F_{11,714} = 24.76$) images ($p < 0.01$). However, we did not find a similar correlation between the sense of depth and the contrast of the darker parts of the images.

The smaller range of highlight contrast in the Leaves image is due to the higher mid-tone level of that image, which limits the available high-end contrast. The mid-tone levels of the Leaves image ranged between 56-70 cd/m^{2}while the mid-tone levels of the Ocean image ranged between 19-34 cd/m^2. In all cases, the background level was set to 0, resulting in luminance levels between 0.1-0.2 cd/m^2.

The significance of the highlights in conveying depth is consistent with the observations of [Berbaum et al. 1983], as well as [Meylan et al. 2006] who found that bright specular highlights lead to a "more natural impression." Meylan et al. also pointed out that the strength of this effect can vary between images, which is also con-

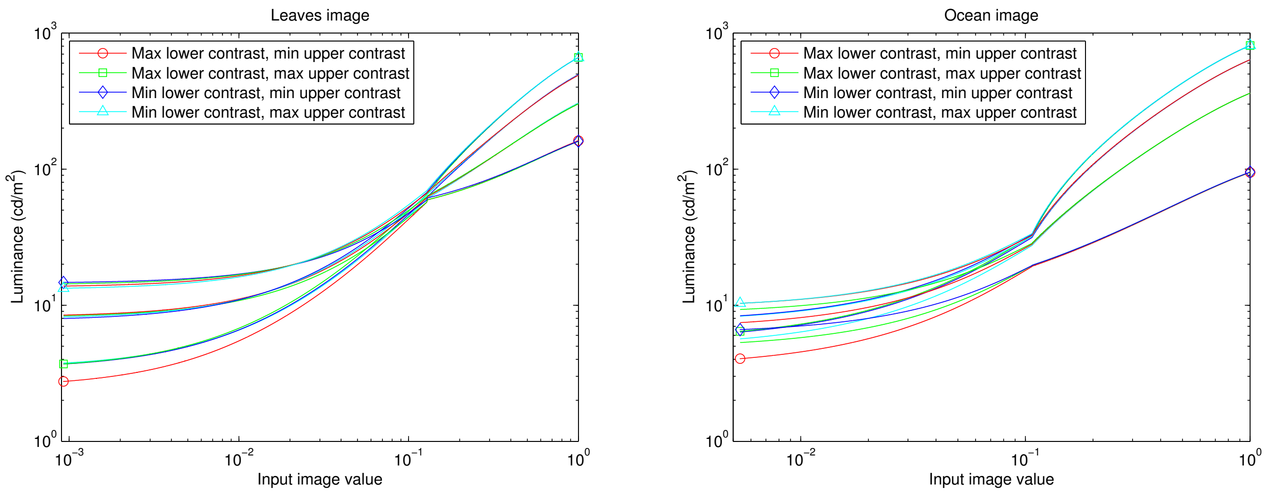

Figure 8: *Logarithmic plots of the tone curves for the two images used in Experiment 3.*

sistent with our results. These observations relate back to Gabriel Lippmann's challenge that a photographic print might someday appear as a window into the world [Lippmann 1908]. If the evolution of photography and the rendering of natural images is toward making it appear as if one were looking out a window, the sensations of both depth and a natural impression should be heightened as we make forward progress. Since the contrast of highlights conveys both, the use of those along with higher-contrast displays brings us closer to our goal.

4 Discussion

The results of our experiments show that the effect of depth within a scene can be heightened by increasing the contrast, both of small-scale features as well as large-scale features, within the scene. Experiments 1 and 2 demonstrated the effect of area contrast on depth and showed how HDR displays can enhance the depth effect even when only boosting area contrast. Experiment 3 then showed that even when area contrast is set to already-high levels such as 200:1 to 500:1 against a black background, it is still possible to obtain an even stronger sense of depth by boosting the contrast of the highlights.

Multiplying the highlight contrast with the mid-tone area contrast, we obtain total contrast ratios of up to 5000:1 in our Experiment 3. Even at those contrast levels, we were able to obtain noticeable differences in subjects' perception of depth. By taking advantage of the full range of available contrast in current and future generations of high-contrast display devices, we should be able to improve viewers' sensations of depth still further, particularly when the depth information from contrast is in line with other depth cues that are present in images. On the other hand, if the depth from contrast conflicts with the other depth cues, high-contrast display devices may diminish rather than enhance viewing experiences, just as inconsistent brightness was observed by [Miles 1953] to diminish binocular viewing experiences.

5 Conclusion

We conducted three experiments in which we observed that increases in contrast resulted in increases in perceived depth in a displayed scene. In our first experiment, we confirmed the findings of Ichihara et al. who had previously explored this relationship with conventional (LDR) displays using random-dot disks. We also expanded upon their analysis by showing that perceived depth also increases with area contrast, and that texture contrast affects that relationship, just as area contrast affects the relationship between texture contrast and perceived depth. And we extended their experiment to show that the same relationships hold when using natural textures instead of random-dot disks.

In our second experiment, we extended the study to the high-contrast range available on HDR displays, using both random-dot disks and natural textures. We confirmed that the same relationships between contrast and perceived depth continue to hold, and that the higher contrast capabilities of HDR displays allow for stronger sensations of depth.

In our third experiment, we showed that the contrast between the highlights and the mid-tones in natural scenes is a strong determinant of the perceived depth within those scenes. Together, our experiments show the depth effect that can come from high-contrast imagery, and that the contrast of both large-scale features and small-scale features contribute significantly to that overall effect.

The current surge of interest in depth perception in visual imagery, at the same time as 3D representations of theatrical feature films are also finding renewed popularity, indicates motivation for further research in this area. In the future, we plan to conduct further studies to examine the relative strengths of different cues, and better quantify the degree to which HDR imagery can have an impact on depth perception.

References

ANDREAS, 2009. Suede leaf. http://www.dreamstime.com/suede-leaf-imagefree191102, accessed Dec. 17, 2009.

BERBAUM, K., THARP, D., AND MROCZEK, K. 1983. Depth perception of surfaces in pictures: looking for conventions of depiction in pandora's box. *Perception 12*, 5–20.

BORING, E. G. 1942. *Sensation and perception in the history of experimental psychology*. D. Appleton-Century company.

BRAINARD, D. H. 1997. The psychophysics toolbox. *Spatial Vision 10*, 433–436.

CIPILOGLU, Z., BULBUL, A., AND CAPIN, T. 2010. A framework for enhancing depth perception in computer graphics. *Proc. 6th ACM Symp. Applied Perception in Graphics and Visualization (APGV)*, 141–148.

FATTAL, R. 2008. Single image dehazing. *ACM Trans. Graph. 27*, 3, Article 72 (9 pages).

GILCHRIST, A. L. 1980. When does perceived lightness depend on perceived spatial arrangement? *Perception & Psychophysics 28*, 6, 527–538.

HELD, R., COOPER, E., O'BRIEN, J., AND BANKS, M. 2010. Using blur to affect perceived distance and size. *ACM Trans. Graph. 29*, 2, Article 19 (16 pages).

ICHIHARA, S., KITAGAWA, N., AND AKUTSU, H. 2007. Contrast and depth perception: Effects of texture contrast and area contrast. *Perception 36*, 686–695.

LIPPMANN, M. G. 1908. Reversible prints. *Communication at the French Society of Physics: Session of March 20, 1908, J. Phys. 4*, 7, 821–825.

MANTIUK, R., AND SEIDEL, H.-P. 2008. Modeling a generic tone-mapping operator. In *Computer Graphics Forum (Proc. Eurographics)*, 699–708.

MEYLAN, L., DALY, S., AND SÜSSTRUNK, S. 2006. The reproduction of specular highlights on high dynamic range displays. In *Proc. of the 14th Color Imaging Conference*.

MILES, P. W. 1953. Anomalous binocular depth perception due to unequal image brightness. *AMA Archives of Ophthalmology 50*, 4, 475–478.

O'SHEA, R., BLACKBURN, S., AND ONO, H. 1994. Contrast as a depth cue. *Vision Research 34*, 12, 1595–1604.

REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph. (Proc. of ACM Siggraph) 21*, 3, 267–276.

RICHTER, J. P., RICHTER, I. A., AND BELL, R. C. 1939. *The literary works of Leonardo da Vinci*. Oxford university press.

ROGERS, B. J., AND GRAHAM, M. E. 1983. Anisotropies in the perception of three-dimensional surfaces. *Science 221*, 4816, 1409–1411.

ROHALY, A. M., AND WILSON, H. R. 1999. The effects of contrast on perceived depth and depth discrimination. *Vision Research 39*, 9–18.

SCHIRILLO, J., REEVES, A., AND AREND, L. 1990. Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information. *Perception & Psychophysics 48*, 1, 82–90.

SEETZEN, H., WHITEHEAD, L., AND WARD, G. 2003. A high dynamic range display using low and high resolution modulators. In *Society for Information Display Internatiational Symposium Digest of Technical Papers*, 1450–1453.

SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Trans. Graph. (Proc. ACM Siggraph) 23*, 3, 760–768.

SEETZEN, H., 2007. Personal communication. At trade shows, HDR images were said by observers to look more "plastic" or "three-dimensional". This was true both for experts as well as a naive audience. However, the effect was found to be image-dependent. In particular, images adapted from LDR content often did not exhibit this effect as strongly.

TRIESMAN, A. 1962. Binocular rivalry and stereoscopic depth perception. *The Quarterly Journal of Experimental Psychology 14*, 1, 23–37.

TROSCIANKO, T., MONTAGNON, R., CLERC, J. L., MALBERT, E., AND CHANTEAU, P.-L. 1991. The role of colour as a monocular depth cue. *Vision Research 31*, 11, 1923–1930.

VON HELMHOLTZ, H. 1924-25. *Physiological Optics / translated from 3d German ed. originally published in 1866*. The Optical Society of America.

WATT, S. J., AKELEY, K., ERNST, M. O., AND BANKS, M. S. 2005. Focus cues affect perceived depth. *Journal of Vision 5*, 834–862.

WIJNTJES, M. W. A., AND PONT, S. C. 2010. Pointing in pictorial space: Quantifying the perceived relative depth structure in mono and stereo images of natural scenes. *ACM Trans. Appl. Percept. 7*, 4, Article 24 (8 pages).