Resolution-Dependence of Perceived Contrast of Textures

Raph Levien

artofcode LLC, 940 Tyler St. Studio 6, Benicia CA 94510

ABSTRACT

Spurred by technological improvements, displays (as well as printers) are increasingly available in a wide range of resolutions. Increased resolution improves perceptual quality in at least two different ways: reducing the perceived contrast of undesirable artifacts (such as halftoning or dithering textures), and increasing the perceived contrast of desirable image features (particularly when rendering text and high precision graphics). Much of the past literature addresses questions of how to optimize one or both of these for a given resolution, but there is little guidance on tradeoffs when the resolution itself is variable.

In this paper, we present an analytic framework for quantifying how the perceived visual contrast of textures changes with resolution, and a simple, tractable model that accurately predicts visual contrast of grayscale-rendered text at different resolutions. These contrast metrics provide a solid basis for evaluating the effectiveness of grid-fitting and similar techniques for perceptually tuned grayscale font rendering, and can also be a useful tool for evaluating engineering tradeoffs such as choosing an optimum resolution relative to cost, speed, or bandwidth constraints.

1. INTRODUCTION

The role of *contrast* is central in visual processing of computer displays. Given a choice, we strongly prefer images of higher contrast. This preference motivates the quest for ever higher contrast ratios in displays, as well as greater ink densities in printing. Further, studies show that contrast is an important predictor of text readability.[1] This paper presents several analytical models and measurement techniques for evaluating the contrast of textures. Section 2 presents a simple analytical model for predicting RMS contrast of text, based on asymmetric pulse-shaped gratings. Section 3 discusses the effects of the contrast sensitivity function (CSF) of the human visual system, especially in the rendering of text. In Section 4, we revisit the perennial question: just how much resolution is enough? Section 5 applies the same contrast sensitivity metric developed for evaluating text quality to halftone patterns, thus demonstrating the usefulness of the visually weighted RMS contrast metric as a unified tool for evaluating display quality in a variety of contexts.

2. CONTRAST OF ANTIALIASED TEXT

The literature contains many different metrics for evaluating contrast. One of the most common is the Michelson contrast definition:

$$C_M = \frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}} \tag{1}$$

Where $L_{\max}$ and $L_{\min}$ are the maximum and minimum luminances.

However, we find this measure to be unsatisfying for evaluating the contrast of antialiased (grayscale) text, as it fails to capture the difference in perceived contrast between high-contrast bilevel text rendering and lower contrast antialiased text. See Figure 1, for example. The Michelson contrast is identical for the two samples, even though the antialiased version clearly has lower perceived contrast.

Instead, we prefer the root-mean-square contrast metric, which is equivalent to the standard deviation of luminance:

Figure 1. Antialiased vs. bilevel display of text.

$$C_{\text{rms}} = \frac{1}{N}\left(N\sum L^2_{(x,y)} - \left(\sum L_{(x,y)}\right)^2\right)^{\frac{1}{2}} \tag{2}$$

In a study of the contrast of natural images, the RMS contrast was found to be the most reliable indicator of the visibility of images.[2]

We first consider renderings of a grating as a simplified model of texture contrast, under two assumptions: that the eye is perfect, and second, with a realistic model of the eye's contrast sensitivity function. The former case is quite straightforward to analyze, and is accurate when the resolution of the display is relatively low.

2.1. Contrast of displayed gratings

Gratings are an important tool for evaluating contrast in imaging systems and the human visual system. In this section, we analyze the relative contrast of gratings when displayed at various resolutions. To most accurately model text, our gratings have a square pulse shape (as opposed to sinusoidal), and are in general asymmetrical, meaning that the width of white and black stripes are not equal.

In the following, we consider a grating of period d pixels, with a duty cycle (ratio of white stripe width to grating period) g. When displayed at infinitely high resolution, the RMS contrast is:

$$C_{\text{rms}}(g) = \sqrt{g(1-g)} \tag{3}$$

Note that this contrast figure is independent of the period of the grating.

Now consider the rendering of this grating at finite resolution, using a simple box filter to model antialiasing. To simplify the analysis, assume that the stripes are at least one pixel wide, i.e. $1 \leq dg \leq d-1$.

Examining the rendering of such a grating, we see that, in an N pixel wide cross section, there are $N_{\text{wb}} = N/d$ white-to-black transitions, and an equal number $N_{\text{bw}} = N/d$ of black-to-white transitions. The rendered pixel at each such transition has a value uniformly distributed between 0 and 1. Thus, the expected value of $L_{(x,y)}$ is $1/2$, and the expected value of $L^2_{(x,y)}$ is $1/3$.

The remaining pixels are all either white or black. The relative number of white pixels N_{w} and black pixels N_{b} is determined by the fact that the expected value for pixel luminance is g:

$$\begin{aligned}
N &= N_{\text{bw}} + N_{\text{wb}} + N_{\text{b}} + N_{\text{w}} \\
N_{\text{bw}} &= N_{\text{wb}} = N/d \\
0.5 \cdot N_{\text{bw}} + 0.5 \cdot N_{\text{wb}} + 1 \cdot N_{\text{w}} + 0 \cdot N_b &= gN \\
N_{\text{w}} &= (g - 1/d)N
\end{aligned} \tag{4}$$

Plugging these values into Equation 3, taking the limit as N tends to infinity, and simplifying yields:

$$C_{\text{rms}}(g,d) = \sqrt{g(1-g) - \frac{1}{3d}}, \quad \text{when } 1 \leq dg \leq d-1 \tag{5}$$

Now, we come to a key result. The relative contrast of the grating rendered at finite resolution, compared with an infinitely high resolution rendering, is:

$$R_{\text{rms}}(g,d) = \frac{C_{\text{rms}}(g,d)}{C_{\text{rms}}(g)} = \sqrt{1 - \frac{1}{3dg(1-g)}}, \quad \text{when } 1 \leq dg \leq d-1 \tag{6}$$
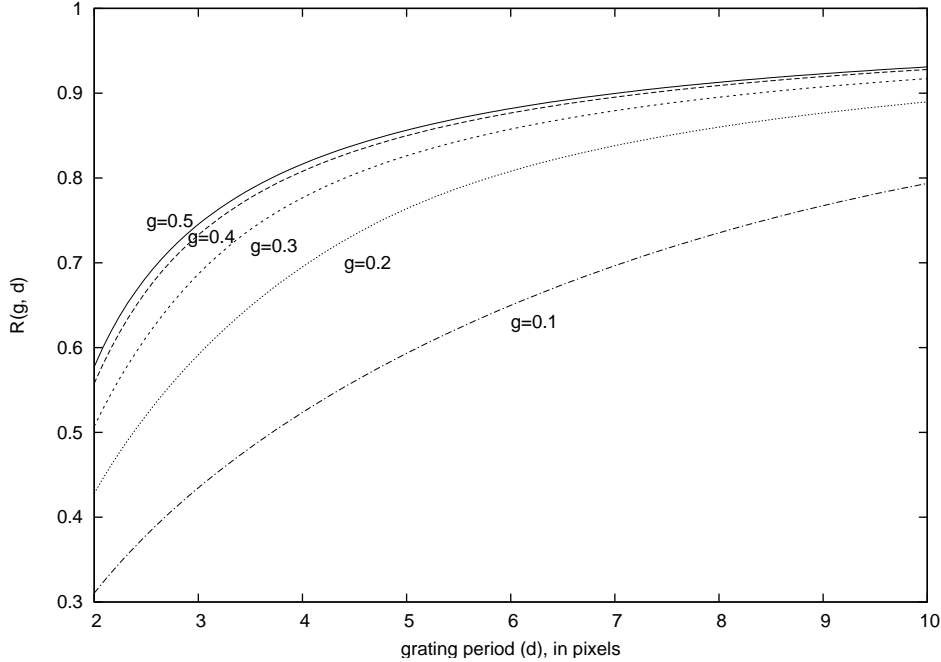
Figure 2. Relative RMS contrast of gratings as a function of period.

2.2. Thin stripes

A similar analysis applies when the stripes are thinner than one pixel. Because of white/black symmetry, we can assume $dg \leq 1$ without loss of generality. In this case, for a band N pixels wide, the population of pixels is:

$$
\begin{aligned}
N_{\mathrm{b}} &= N(1 - (1/d)(1 + gd)) \\
N_{\mathrm{w}} &= 0 \\
N_{\mathrm{bw}} &= N_{\mathrm{wb}} = gN \\
N_{\mathrm{g}} &= N_{\mathrm{b}} - N_{\mathrm{wb}} - N_{\mathrm{bw}} = N(1 - (1/d)(1 + gd) - 2g)
\end{aligned}
\tag{7}
$$

Of these pixels, edge pixels N_{wb} and N_{bw} are uniformly distributed between 0 and dg, while gray pixels N_{g} have the constant value dg. Plugging all this in to Equation 3, and taking the limit as $N \to \infty$, we derive:

$$
R_{\mathrm{rms}}(g, d) = \sqrt{\frac{g}{1 - g}\left(d\left(\frac{1 - gd}{3}\right) - 1\right)}, \text{ when } gd \leq 1
\tag{8}
$$

The RMS contrast for rendered gratings, covering both the cases of Equation 6 and Equation 8 is shown in Figure 2:

2.3. Contrast of displayed text

Gratings are relatively easy to analyze, but are not a real application of displays. However, people do spend a lot of time reading text. Fortunately, we can predict the relative contrast of text as a function of text size using these gratings as a model.

This model has two parameters, which vary depending on the choice of font. One parameter, g, is simply the duty cycle of the grating. The other, β, represents the ratio of grating period to the "em size" of the font. We then predict the relative contrast of a text display of y points per em (ppem) using this simple formula:
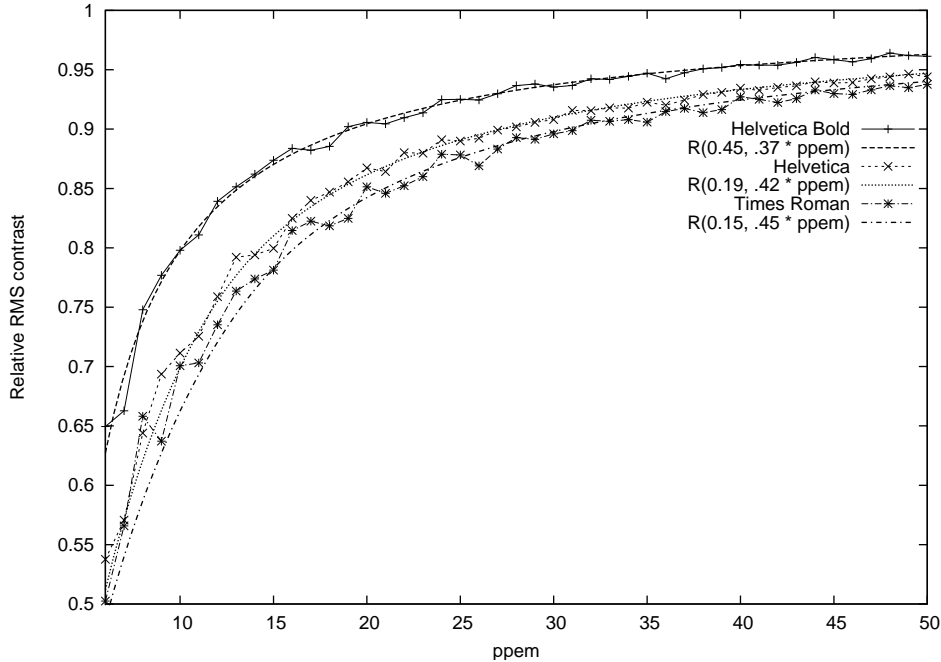
Figure 3. Relative RMS contrast of text as a function of size (pixels per em).

$$R = R_{\mathrm{rms}}(g, \beta y) \tag{9}$$

Figure 3 demonstrates the accuracy of this prediction, based on parameters empirically chosen to fit the model for three fonts. Times Roman is represented by $g = .15, \beta = .45$, Helvetica by $g = .19, \beta = .42$, and Helvetica Bold by $g = .45, \beta = .37$. Unsurprisingly, bolder fonts are represented using duty cycles closer to 0.5. Note that our contrast measure is symmetric with respect to white and black polarity, justifying the convention of a smaller duty cycle for a thinner font, whether the actual display is black-on-white text or the reverse.
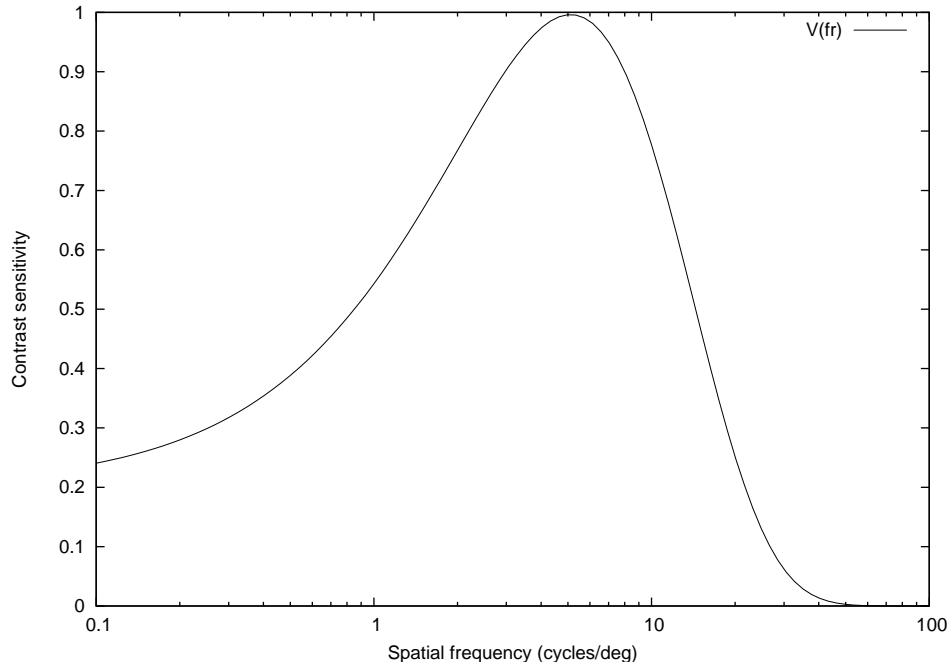
3. CONTRAST SENSITIVITY

The analysis of the previous section is only valid under the assumption that the human visual system is equally sensitive to contrast at all spatial frequencies rendered by the display. This assumption is roughly valid at low resolutions (when the eye can perceive individual pixels), but not as the resolution scales up.

The eye has a peak contrast sensitivity at around 5 cycles/degree, and falls off rapidly at higher spatial frequencies. One simple, widely used model for contrast sensitivity is the Nill-Bouzas model[4]:

$$V(f_r) = (0.2 + 0.45f_r)\exp(-0.18f_r) \tag{10}$$

Where f_r is the spatial frequency in units of cycles per degree. Figure 4 shows the plot of this contrast sensitivity function.

We incorporate this contrast sensitivity function straightforwardly into our contrast metric: to compute the spectrally-weighted RMS contrast of an image, we first compute the spectrum of the image, then multiply by the contrast sensitivity function (rotating around the origin for radial symmetry), and finally compute the RMS of the resulting weighted spectrum. This formulation is equivalent to convolving the image with a lowpass filter with a spectrum equal to the contrast sensitivity function, then computing the RMS of the blurred image

Figure 4. Nill-Bouzas contrast sensitivity function.

The inclusion of the contrast sensitivity function has a fairly subtle effect on the measured contrast of antialiased text, as can be seen in Figure 5. The most significant effect is a more rapid approach to the asymptote of 1.0. At a reading distance of 35cm, 10 point Helvetica type reaches a relative contrast of 0.95 at about 170 dpi, or 24 ppem. Assuming a perfect eye, achieving the same relative contrast would require about 400 dpi, or 56 ppem. The energy content of the difference between the 170 dpi display and an infinite-resolution one explains the difference; most of this energy is at high spatial frequencies that can be perceived only dimly.

4. HOW MUCH RESOLUTION IS ENOUGH?

In his Ph.D. thesis,[5] Avi Naiman posed the question: how much resolution is enough? He gives the figure of 580 pixels per inch as something of an upper bound—it should not be possible for the human visual system to perceive differences above this resolution at all. However, we would like to address a slightly different question: how much resolution is enough for practical purposes?

Given that one of the most practical purposes is display of text at normal reading size, for use at normal reading distance, we frame this question more precisely: how much resolution is required to display 10 point text at 95% of the perceived contrast of an infinite-resolution display? The answer is 170 dpi. For the less ambitious goal of 90% perceived contrast, the required resolution is 115 dpi.

4.1. Reading distance

The assumption of a normal reading distance of 35–40cm is probably valid for most desktop and laptop computer displays, but is likely too generous for most uses of smaller devices such as PDA's and cell phones. For example, the present author finds that comfortable reading distance for his PDA and cellphone are about 25cm, as opposed to 35cm for his laptop. Since the device resolution required to maintain a constant angular resolution scales inversely with reading distance, 95% perceived contrast requires about 240 dpi.

The comfortable reading distance for small devices depends largely on *accommodation,* the ability of the eye to focus at a range of distances, in turn largely a function of age.[7] Children can accommodate up to 14 diopters, allowing reading distances as small as 7cm. The amplitude of accommodation then decreases by roughly 0.16
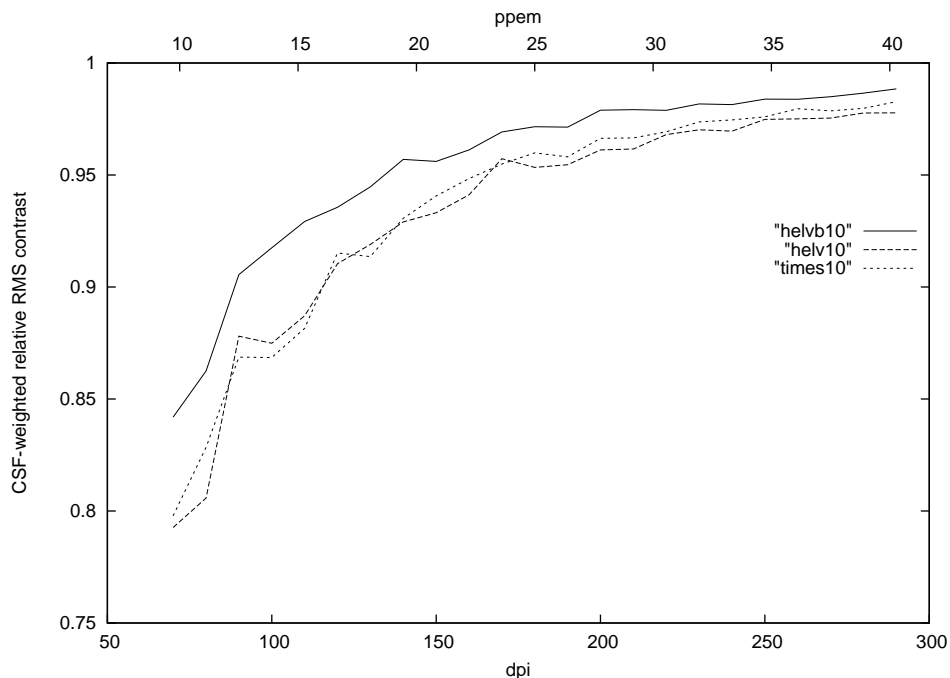
Figure 5. Relative contrast as a function of resolution; 10 point type, 35cm reading distance.

diopters per year until it levels out at around 2 diopters, at which point glasses are necessary for comfortable reading.

Given the increasing ability of portable devices to display content such as Web pages, and the limited display size required by portability constraints, it makes sense that youthful users will gravitate towards greater amounts of information on tiny displays, and a correspondingly smaller reading distance. The exact distance (and hence resolution) will depend on the individual user, but resolutions of up to 800 dpi will prove necessary to deliver the full experience of high-contrast text.

4.2. Grid fitting

While many antialiased text rendering systems use simple box-filter subsampling to achieve smooth curves and minimally distorted letterforms, others use some type of *grid fitting* technique to improve contrast. In general, grid fitting techniques perform some distortion on the glyph outlines in order to align edges to pixel boundaries. One common technique is to apply a hinting algorithm originally intended for bilevel rendering (such as TrueType or Type 1 hints) to the outline before grayscale rendering. Other, more sophisticated approaches optimize directly for antialiased rendering.[6]

In most cases, such techniques force a tradeoff between contrast enhancement and letterform distortion. The contrast metrics presented in this paper provide a quantitative basis for evaluating exactly how much contrast is gained. A metric for the effects of distortion on readability and percieved quality is beyond the scope of this work.

Since grid fitting increases contrast, it allows a lower resolution to suffice, if all that is measured is contrast. Assuming that the distortion of letterforms is tolerable, it relieves some of the pressure on the display technology to increase the resolution. The evaluation of existing and novel grid fitting techniques using the contrast framework remains a problem for future research.
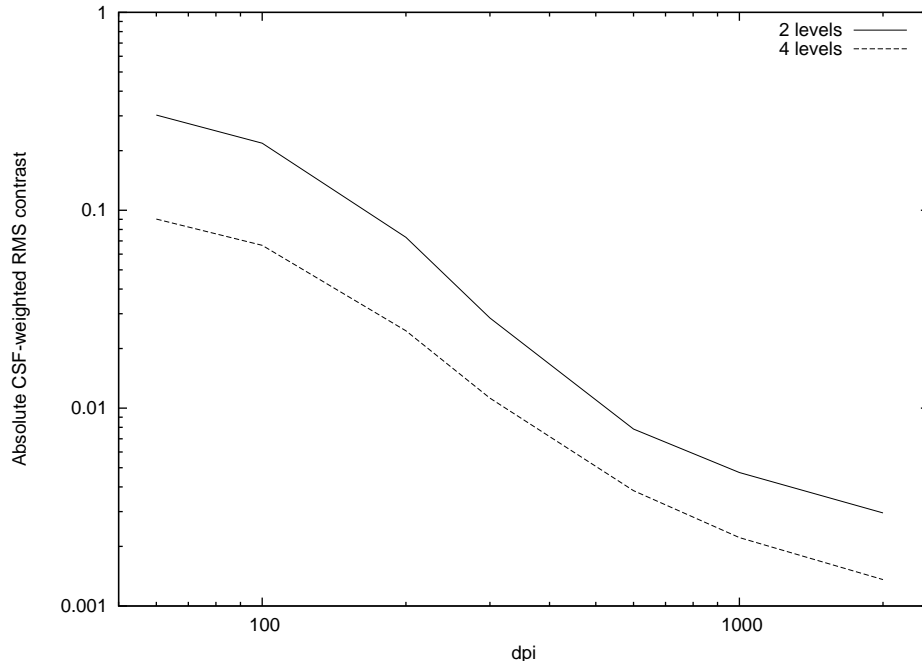
Figure 6. Absolute CSF-weighted RMS contrast as a function of resolution; Even Tone Screening; 35cm

5. CONTRAST OF HALFTONES

The weighted RMS constrast metric is familiar in the context of design and evaluation of halftone patterns.[8] In this sense, it is usually referred to as an *error metric*, as it represents the perceived contrast of the error between the original continuous tone image and its halftoned rendition. Where the goal in rendering text is to maximize texture contrast, in halftoning the goal is to minimize it, while preserving the contrast of the original image unmodified.

Overall contrast is affected by both the resolution and the number of levels of quantization.[9] Figure 6 shows the result of halftoning the standard Lena image with Even Tone Screening,[10] using 2 and 4 levels. For example, achieving a dither pattern contrast of 0.004 (roughly equivalent to quantization noise at 8 bits per pixel) requires 1200 dpi with 2 levels, but only 580 dpi with 4 levels, providing strong justification for the popularity of 2 bit-per-pixel variable dot inkjet printers today,

These results are based on the assumption of an ideal device, both in terms of lack of dot gain and linearity of the actual dot sizes with respect to the dot size codes sent to the printer. Applying the contrast metrics to a real device should take into account these real-world imperfections as well.

6. DISCUSSION

We have presented a framework for evaluating the contrast of both desirable (text) and undesirable (dither pattern) textures, using an RMS contrast metric weighted with a contrast sensitivity function. In particular, this metric clearly highlights the tradeoffs between device resolution and perceived quality.

Gratings are particularly susceptible to analysis. We have shown how to to easily compute unweighted RMS contrast for a class of asymmetric pulse-shaped gratings. Further, we have shown that these gradients are a very close fit to RMS contrast of text rendered using scalable fonts.

A spectrally weighted RMS contrast metric, using a model of the human visual system's contrast sensitivity function, provides an even more accurate model of perceived quality than the unweighted version, but is of course more difficult to compute. We have shown that, for text, the effect of adding the spectral weighting is subtle,

having generally the same shape and characteristics as the unweighted case. However, convergence to the ideal relative contrast of 1.0 occurs more rapidly as resolution scales upward. Using this metric, we conclude that 170 dpi displays are sufficient to display very high contrast (95% relative to the infinite resolution case) text under normal viewing conditions.

The same metric is equally useful for analyzing unwanted textures, in particular halftone dither patterns. It gives guidance in the tradeoff between resolution and quantization coarseness. Using this metric, we conclude that an ideal 1200 bilevel printer is adequate to produce very low contrast (equivalent to 8 bits per pixel image quantization) dither patterns, and that a variable-dot printer capable of 2 bits per pixel produces similar texture contrast at 580 dpi, a bandwidth savings of more than 2:1.

We propse that the weighted RMS contrast metric is a generally useful tool for evaluating display quality, both in familiar contexts such as halftoning, and potentially new applications such as font rendering. We hope that future researchers will make good use of the metric in evaluating the quality of imaging techniques, and provide useful technological guidance for delivering the highest possible image quality, especially in constrained environments.

REFERENCES

1. Gordon E. Legge, Gary S. Rubin and Andrew Luebker. *Psychophysics of Reading: V. The Role of Contrast in Normal Vision.* Vision Research, 27(7): 1165–1177, 1987.
2. Bex, P. J., and Makous, W. (2002). *Spatial frequency, phase and the contrast of natural images.* Journal of the Optical Society of America A, 19(6): 1096–1106.
3. Gordon E. Legge, Denis G. Pelli, Gary S. Rubin and Mary M. Schleske. *Psychophysics of Reading: I. Normal Vision.* Vision Research 24(2): 239–252, 1985.
4. N. B. Nill and B. H. Bouzas. *Objective image quality measure derived from digital image power spectra.* Opt. Eng 31: 813–825, 1992.
5. Avi Naiman. *The Use of Grayscale for Improved Character Presentation.* Ph.D. Thesis, U. Toronto, CSRI-253, 1991.
6. Kevin O'Regan, Nicole Bismuth, Roger D. Hersch, and Alexandros Pappas. *Legibility of Perceptually-Tuned Grayscale Fonts.* Proc. ICIP-96, 1: 537–540, ed. P. Delogne, 1996.
7. C. Miege. *Age-related changes in accommodation and refraction for presbyopes.* Insight Professional J, 1997.
8. Henry Kang. *Digital Color Halftoning.* SPIE Press, 1999.
9. Jennifer Gille, Martin, R. and Larimer, J. *Spatial resolution, grayscale, and error diffusion trade-offs: Impact on display system design.* Conference Record of the 1994 International Display Research Conference, 381–385, 1994.
10. Raph Levien. *Practical Issues in Color Inkjet Halftoning.* Color Imaging VIII: Processing, Hardcopy, and Applications, 5008-57, 2003.