Published in final edited form as:

J Soc Inf Disp. 2012 November 1; 20(11): 616–623. doi:10.1002/jsid.127.

Electronic magnification and perceived contrast of video

Andrew Haun, Russell L Woods, and Eli Peli

Schepens Eye Research Institute, Mass Eye and Ear, Harvard Medical School, Boston MA

Abstract

Electronic magnification of an image results in a decrease in its perceived contrast. The decrease in perceived contrast could be due to a perceived blur or to limited sampling of the range of contrasts in the original image. We measured the effect on perceived contrast of magnification in two contexts: either a small video was enlarged to fill a larger area, or a portion of a larger video was enlarged to fill the same area as the original. Subjects attenuated the source video contrast to match the perceived contrast of the magnified videos, with the effect increasing with magnification and decreasing with viewing distance. These effects are consistent with expectations based on both the contrast statistics of natural images and the contrast sensitivity of the human visual system. We demonstrate that local regions within videos usually have lower physical contrast than the whole, and that this difference accounts for a minor part of the perceived differences. Instead, visibility of ‘missing content’ (blur) in a video is misinterpreted as a decrease in contrast. We detail how the effects of magnification on perceived contrast can be measured while avoiding confounding factors.

Keywords

video scaling; interpolation; digital zoom; super-resolution

1. Introduction

Magnification of digital imagery results in a decrease in angular resolution, and so the resulting image is often perceived as blurred. The perceptual impact of digital magnification is not well understood except in that blur and interpolation artifacts are objectionable, but the goal of improving ‘super-resolution’ algorithms is nonetheless to produce magnified digital images with minimal impact on perceived quality of the result (Farsiu et al 2004¹). While developing visual rehabilitation aids that use electronic magnification (Goldstein et al 2003²; Peli et al 2009³), we have noticed an apparent attenuation of image contrast with magnification, and reports indicate that the effect has been noted in other contexts (e.g. Knoche et al 2007⁴). Here, we measure this effect using motion video and attribute its cause both to physical variations in local contrast within natural images, and to a perceptual effect linked to the visible resolution limit of the magnified videos.

In this paper we are interested in the effects of magnification on *perceived luminance contrast* of an image. Luminance contrast is a basic statistic of any image, but for complex images (and even for simple patterns; Peli 1997⁵) contrast is difficult to summarize with either physical or perceptual measures (Peli 1991⁶). Because perceived contrast is such an important feature of image quality, it is typically included in the early computations of many image quality metrics (cf. Wang & Bovik 2005⁷), but the highly nonlinear computations underlying perceived contrast of complex images are still not well-understood. Reduction of a complex image’s contrast makes it look faded or washed out – anything less than ‘true’ contrast is seen as a decrement in image quality, and in this sense subjects do seem to be able to identify the global contrast of a real-world scene (e.g. as in Bex & Makous 2002⁸)

independent of other properties such as sharpness (May & Georgeson 2007⁹). On this point it is necessary to differentiate not only between physical and perceptual contrast, but also between physical and perceptual blur. To blur an image is to remove detail from it, and this usually also involves loss of contrast – this is why blur and contrast attenuation are easily confounded especially by naive observers (Watt & Morgan 1983¹⁰; Roufs 1989¹¹). However, digitally magnified images are not physically blurred – no detail is removed, and the luminance distribution is unchanged. Instead, magnification reveals the image's limit of encoded detail which in an unmagnified image may have been invisible. So, magnified images may *appear* blurred relative to unmagnified originals, although they are structurally the same apart from their scale.

Two hypotheses are available to explain any difference in perceived contrast between normally displayed and magnified video. First, as discussed above, contrast is perceptually related with sharpness and blur, so it may be that when blur or pixelation is visible in an image, it provokes a sensation of overall contrast loss. That is, the explanation may be entirely based in perception. Alternatively, there may be a mundane explanation: because of the heterogeneity of natural image structure, any subregion of an image containing varied content is likely to have a smaller range of luminance and contrasts than the full image, so that observed differences in contrast may have an entirely physical basis. This second hypothesis can only hold if, when estimating video contrast, our subjects are estimating the physical distribution of luminances in the stimuli.

We evaluated the effect of electronic magnification on perceived contrast by having subjects equate the perceived contrast of a normally-displayed video clip with a magnified version of the same. We carried out our experiments using bilinear interpolation, treating it as the most basic plausible electronic magnification algorithm. We found that the difference in perceived contrast was strongly affected by the presence of content outside the magnified area, supporting the physical contrast difference hypothesis, but that there was an additional, purely perceptual effect not explained by the presence of the cropped content. To test the second hypothesis that this effect was due to the perceived blur of the magnified videos, we repeated the experiment at multiple viewing distances, including a distance great enough that the pixelation or blur of the magnified video should have become invisible (Peli 2001¹²). Under this condition, we found that the perceived contrast effect was nearly abolished when cropped original videos were used, but it remained when full-sized originals were used, lending further support to the notion that while the perceived contrast is related to blur it is also affected to an extent by a physical difference in local versus global contrast. Also, we found that simultaneous comparisons of the contrast of videos of different angular size presented against a blank background can be severely confounded with video size, and demonstrate a set of controls that make good, objective comparisons possible.

2. Methods

2.1 Stimuli

Stimuli were 100 3-second video clips drawn at random (excluding segments containing scene cuts) from two DVD movies and played continuously in a forward-backward loop until subject response. All videos were displayed in grayscale as uncompressed (post extraction from the DVD) avi files – i.e. playback did not involve any decoding or decompression. We assumed that the videos were intended for display on an ordinary display with a gamma of ~2.0, but we wanted our videos to look 'normal' despite being viewed on a linearized display (see 2.4 *Equipment* below), so we 'undid' the video gamma compression by raising video pixel intensities to a power of 2 (cf. Bex et al 2005¹³). As illustrated in Figure 1, videos could be displayed in one of three ways: 1) 'full-size', where each frame was 360×360 pixels taken directly from the center of a DVD frame, 2)

‘cropped’, where only the central portion to be magnified on a given trial was displayed (e.g. a 120x120 pixel video for 3x magnification) at its original resolution, and 3) magnified, where the central portion of the original video was expanded through bilinear interpolation to 360x360 pixels. For full-sized and cropped videos, video pixels and monitor pixels were the same size, but for magnified videos, video pixels increased in size with magnification. Videos were centered 184 monitor pixels to the left and right of the center of the 800 by 600-pixel display. Background (non-video surround) pixels were set according to the experiment condition.

2.2 Procedure

Subjects performed a discrimination task, choosing which of two side-by-side videos seemed to have the higher contrast. Two definitions for ‘higher contrast’ were given to each subject: ‘larger range of grayscale values’ (for the subjects more familiar with psychophysics or image processing) and ‘brighter whites/brights and darker blacks/darks’ (for both experienced and naive subjects). The experimenter explicitly stated that the subjects were not to judge *sharpness*, since on every trial the unmagnified (original) video would obviously be sharper than the magnified video. Subjects were also instructed to choose not on the basis of single local features within videos, but rather to try to estimate the overall contrast of the video over both spatial extent and over time (admittedly difficult and subjective, and we do not doubt that subjects varied in their ability to accomplish this. A 1-up 1-down staircase adjusted the contrast of the original video (full-size or cropped) in 0.05 log unit steps from trial to trial, according to whether on the previous trial the original video was chosen as having higher contrast (resulting in a decrease in original video contrast) or whether the magnified video was chosen (resulting in an increase in original video contrast). This procedure adjusted the contrast of the original video to match the perceived contrast of the magnified video. Each separate staircase ran for 60 trials.

Original and magnified video contrast were set by adjusting the entire video’s RMS (root mean square) contrast: $V' = 10^c(V - \mu_V) + \mu_V$, where V is the source video, V' is the adjusted video, c is the contrast change in log units with respect to the original contrast, and μ_V is the mean value of all pixels in the video. To allow original video contrast to be set to physically greater contrast than the magnified video if necessary, magnified video contrast was fixed in each experiment to -0.2 log units below its original contrast (i.e. 63% of original); cf. Bex & Makous 2002 for a similar procedure and rationale.

Separate staircases were used within a block of trials for trials with the original videos on different sides of the display (i.e. each block of trials consisted of interleaved left-side and right-side staircases), and for different magnification levels. At the end of each experiment, trials were binned by original video contrast (left- and right-side staircases were combined) and fit with a logistic function estimating the proportion of trials at a given original video contrast where the original video would be chosen as having higher contrast than the (fixed contrast) magnified video:

$$p(C'_{\text{original}} > C'_{\text{magnified}}) = 1 / (1 + \exp(-(c_{\text{original}} - c_{\text{match}}) / \lambda)) \quad (1)$$

Here C' denotes *perceived* contrast of the videos and c_{test} the *physical* contrast of the original (unmagnified) video on a given trial. The fitted value c_{match} is the physical contrast of an original video that yields a perceptual match between original and magnified video C' values. λ is set by the slope of the function, being proportional to the width of the transition between seeing the original video as higher-contrast and seeing the magnified video as higher-contrast, and therefore can be taken as a measure of the subjects threshold for seeing

a change in the contrast of the original video, although the procedure was not optimal for making good measurements of λ .

2.3 Subjects

Six subjects participated in Experiment 1; four of the six in Experiment 2; and five of the six in Experiment 3. Subjects were in the age range of 22 to 50y, all with normal or corrected-to-normal visual acuity and with no known visual impairments. Subjects viewed the display at 1, 3, or 5m depending on the experiment.

2.4 Equipment

The display used was a Sony Trinitron p1130 CRT, run at 800×600 (0.476 px/mm) resolution and 144Hz vertical retrace (each video frame was displayed six times for a framerate of 24fps). The display luminance/voltage function was linearized by adjusting the gamma of the three color guns via the video hardware (nVidia GeForce 9400GT). Mean display luminance was 47cd/m². Experiments were carried out on a Windows computer system running Matlab 7.5 with the Psychophysics Toolbox extensions (Brainard 1997¹⁴, Pelli 1997¹⁵).

3. Effects of Cropping on Video Contrast

Physical differences in contrast between original (unmagnified) and magnified videos can be evaluated directly if we adopt a measure of contrast. First, we can consider the differences between local – i.e. to-be-magnified – and global scene contrast. We analyzed local versus global contrast by measuring the RMS value of non-overlapping samples of various sizes over all 72 frames across each of the 100 video clips. The largest tile was 360 pixels, constituting the entire area of the frame and yielding a single contrast, the ‘global’ RMS value; the next largest was 180 pixels, half the size of the frame and yielding four contrast samples; the next was 120 pixels, yielding 9 samples; and so on, for a total of 12 scales (360, 180, 120, 90, 72, 60, 45, 40, 36, 30, 24, and 20 pixels).

Figure 2 shows how the contrast of our stimulus videos depended on the area analyzed. Each set of measurements was normalized to each frame’s global RMS contrast before summary statistics were computed, so that global contrast values here take on a value of 1. As region size decreases, RMS contrast also decreases (circular symbols). The relationship is well-described as a decrease in contrast proportional to the square root of the magnification factor (dashed red line).

The decline in contrast is partially due to the 1/f amplitude spectra typical of natural imagery, as decreasing the sample size excludes higher-power low frequency contrasts (large, but gradual, variations in luminance over the image). This is demonstrated by randomizing the phase angles of the fast Fourier transform (FFT) coefficients. We did this by replacing the phase spectrum of each analyzed video frame with the phase spectrum of a frame-size sample of Gaussian noise. In the spatial domain, this results in a random distribution of the frame’s contrasts over its area, without changing their amplitudes. The square symbols in Figure 2 show how contrast decreases with sample size when the phase spectrum of each original (full size) frame has been scrambled, revealing the contribution of the amplitude spectrum to the contrast decline.

The contribution of the phase spectrum is revealed by flattening the amplitude spectrum of the original image without changing the phase. This is accomplished by setting all FFT coefficient magnitudes to a constant value without altering the phase angles. As shown by the triangle symbols in Figure 2, the phase spectrum is also responsible for a part of the decrease in contrast with decreasing sample area, albeit a smaller part than the amplitude

spectrum. This decrease can be understood as due to natural images' relative lack of stationarity: some regions of an image may be thick with contrast structure, while others may be nearly blank.

Figure 2 also shows how contrast changes with sample area when there is no natural structure at all, by applying the same analysis to Gaussian white noise (which have constant coefficient magnitude and random phase). White noise contrast (X symbols) does not decrease with magnification, at least not within the bounds of our measurements (naturally it must decrease eventually as the number of pixels is drastically reduced – the smallest sample here was 20x20 pixels, so 400 pixels are apparently enough to preserve the contrast structure of white noise).

This analysis demonstrates simply that a decrease in contrast with decreasing sample area (or magnification) is a consequence of the natural structure of a digital image of the real world. If this decrease is responsible for what we have observed as the result of magnifying digital images, then it can serve as a prediction for the magnitude of the effect.

We must also confirm that our interpolation method did not itself change the contrast of the videos. A nearest-neighbor interpolation (to an integer magnification) would perfectly preserve the luminance distribution of the original video, but this method is almost never used in modern applications, as it introduces sharp edges and flat, square surfaces which are distractingly unlike any real-world image. The bilinear interpolation method we used produces a smoothed version of the nearest-neighbor method, but the smoothing has next to no effect on the luminance distribution. Figure 3 shows that the luminance distribution of a frame is not significantly affected by linear interpolation to four times its original size. It is smoothed, but its structure, mean, and standard deviation – the RMS contrast – are not noticeably changed.

4. Results

4.1 Experiment 1: Varying Magnification Level with Full-sized and Cropped Originals

The effect of magnification on the difference in perceived contrast between original and magnified videos was tested using a matching procedure at three magnification levels: 2.0x (180px magnified to 360px), 4.0x (90px magnified to 360px), and 6.0x (60px magnified to 360px). Those magnified videos were compared either with full-size (360px) unmagnified (original) videos or with that original video cropped to contain only the content displayed in the magnified video. For the cropped original condition, the two videos on each trial were identical except for scale. The cropped original videos were 180px, 90px or 60px in size for the 2.0, 4.0 and 6.0 magnification levels, respectively. The full-size and cropped original-video conditions were conducted in separate blocks with half of the subjects performing the full-sized original-video condition first. All three magnification levels were assigned their own left- and right-side staircases and interleaved randomly in a block of trials (so there were six staircases). Subjects viewed the videos from 1m, so the 360px videos subtended 9.8° of visual angle. The display background (the part of the display not occupied by the videos) was fixed at the mean display luminance (47cd/m²), except as noted later.

Results are shown in Figure 4 as the ratio of unmagnified RMS matches to the magnified videos. For all subjects, for each magnification level and both original conditions, matches were less than unity – subjects always underestimated the contrast of the magnified videos. Whether the original video was cropped or full-size, the degree of underestimation increased as magnification increased.

Unexpectedly, when magnification was increased beyond a factor of 2, the effect on perceived contrast was greater when the original video was cropped to match the content of the magnified video. In both conditions, the magnified videos were the same; i.e. there is no reason to suspect that changing the size of the original videos (cropping them) should change the perceived contrast of the magnified videos. So, the difference in results of the two conditions must be due to changes in the perceived contrast of the original videos. However, as demonstrated in the previous section, physical video contrast *decreases* as a video is cropped, and thus the subjects should have required less of a decrement in contrast to match original to magnified video contrast. The pattern of results shown in Figure 4 indicates that reducing the size of the original video caused its perceived contrast to *increase*. We suspected that there was some confounding interaction between the size of the cropped original videos and the fixed structure of the display background, which led to our second experiment.

4.2 Experiment 2: Effect of the Display Background on Perceived Contrast

Experiment 2 repeated the conditions of Experiment 1 (both cropped and full size original blocks were run in alternating order), except that now the display background was controlled in one of three ways: the background luminance could be matched, frame for frame, to the average luminance of the video frame on the respective side of the display; the background could be filled with dynamic Gaussian white noise (RMS = 0.1) around display mean luminance; or both manipulations at once (noise plus frame-matched luminance). We reasoned that there were two most likely causes of the size effect on cropped video contrast. First, the mean luminance of our video clips was rarely the same as the display mean, so there was usually contrast between the video DC and the background. Decreasing the size of the video would shift this DC-background contrast towards higher spatial frequencies, where they might account for more of the observer's overall contrast judgment (Bex & Makous 2002⁸, Haun & Peli 2012¹⁷). Matching video mean and background luminances would serve to decrease this effect. For the cropped condition, this meant that the entire display background was set at the same luminance since magnification did not affect the mean value of a frame; but for the full size condition, the original and magnified videos could have slightly different mean luminance on any given frame, so the luminance of each video frame was matched by the luminance on the corresponding half of the screen. Second, there could be contrast-contrast effects, where the lack of surround contrast resulted in a release from suppression of the central display region (Chubb, Sperling, & Solomon 1989¹⁸). By adding dynamic contrast noise to the background we aimed to ameliorate such effects.

Figure 5 shows the results of Experiment 2 and replots the results of Experiment 1 (minus 2 subjects who did not return for experiment 2). The main effect of magnified video contrast underestimation was preserved in all the background conditions, but was greatest in the original condition (fix/blank) (Fig.5a). The greatest reduction in effect size was seen when the background was both luminance matched and filled with noise (vary/noise). When the original video was full-sized, background manipulations had much less influence on the perceived contrast effect (Fig.5b). To reveal the effect of cropping on the perceived contrast of the original (unmagnified) videos, we subtract the full-size original data from the cropped original data (Fig.5c) – since the magnified video contrasts were the same in every condition, their contribution is in this way nullified. On these axes, positive and negative values indicate respectively that cropping increased or decreased perceived contrast of the original videos. In the original condition (solid round symbols), cropping a video to a 1/4 or 1/6 its original size seems to have increased its perceived contrast by about 10%. However, when the background was frame-luminance matched and filled with noise (open triangle symbols), there was an overall decrease in original video perceived contrast as a result of

cropping, with the decrease more or less constant. The other two controls resulted in cropping effects similar to, but less than, the original condition.

4.3 Discussion

The major component of the perceived difference in contrast between original and magnified video is perceptual, not physical. Once the influence of size change (cropping) was excluded, we found that the difference in perceived contrast averaged between 10% and 15% (Fig.5a). Introducing a real physical (RMS) difference between the test stimuli, in the full-size original condition, (Fig.5b), only slightly increased the magnitude of the effect, nowhere near the $1/m$ relationship predicted if subjects were actually matching video RMS contrast.

It is interesting that the structure of the background had such a significant effect on the perceived contrast of our videos, but within the scope of this study we cannot speculate reasonably as to the causes of these effects. Our purpose was to eliminate the effects of cropping on the perceived contrast of the original videos so that we could obtain valid estimates of the contrast of the magnified videos. It appears that the vary/noise condition was best able to cancel the effects of cropping on the contrast of the original videos (Fig.5c), so we proceeded to Experiment 3 with these settings.

4.4 Experiment 3: Varying Viewing Distance

In our last experiment, we addressed the relationship between perceived sharpness and perceived contrast; that is, particularly for naive observers, judgments of high and low contrast are normally related with judgments of sharpness and blur. This is a natural conflation, since 'blur' usually occurs in the transitive sense, as something that is done to an image, i.e. the removal of finer spatial detail. With magnification, however, nothing is removed – rather, the absence of higher spatial frequency contrasts is revealed as the pixels become visible. The likely cause of the decrease in perceived contrast is the now-visible contrast gap at the higher spatial frequencies. If this is so, then closing that gap while maintaining the scale difference between magnified and original videos should reduce the size of the effect.

Subjects viewed the display in separate blocks at distances of 1, 3, or 5m. At each viewing distance, a single magnification (3x) was used, with trials randomly interleaved in two separate staircases in the same procedure as Experiments 1 and 2. Viewing distance order was randomized across subjects. Original videos were full-sized or cropped in separate blocks, as in Experiments 1 and 2. Full size or magnified (360x360px) videos subtended with increasing distance 9.8°, 3.3°, and 2.0°; cropped (120x120px) videos subtended 3.3°, 1.1°, and 0.65°. To avoid the effects of simultaneous contrast between video and background as determined in Experiment 2, the display background was again matched to the frame luminance of the stimulus videos, and filled with dynamic Gaussian noise.

The distances chosen were not arbitrary (Figure 6a). At 1m, the monitor pixels were separated by 1.64 arcmin (minutes of arc), limiting the finest details that could be displayed. Normal human acuity is better than this, with the minimum angle of resolution (MAR) less than 1 arcmin for individuals with better than 20/20 acuity. So, normal observers should have been able to see the finest details in the original resolution videos at 1m, but just barely (if they were presented at high contrast); thus the original videos should have looked as sharp as their digital content allowed (and it is unlikely that any of the videos were of such perfect quality that they would have real detail at the resolution limit). Magnifying a video by 3x effectively increases the video pixel separation by the same factor, so that the finest details in the video will be about 5 arcmin apart, easily discriminable (or just discriminable

to someone with 20/100 acuity); thus the magnified videos will appear blurred. At 3m, the unmagnified video pixels should no longer be discriminable, and for most observers some high frequency contrast will be lost beyond the acuity limit, but the videos would not look blurred; in fact, according to Heinrich and Bach (2010¹⁹), they should look more ‘detailed’ than they did at 1m. The magnified videos at 3m should have video pixels discriminable to the same degree as the original videos at 1m; i.e. they should look as sharp and as ‘detailed’ as the original videos did at 1m. At 5m, neither magnified nor original videos should have discriminable video pixels (except to the most eagle-eyed observer); thus neither should have appeared blurred.

Figure 6b shows that with increasing viewing distance, the difference in perceived video contrast decreases for both comparison conditions (ANOVA main effect of distance, $F_{2,8} = 11.3$, $p = 0.005$), with a greater effect of viewing distance on the cropped-original condition (condition/distance interaction, $F_{2,8} = 7.59$, $p = 0.014$). At 5m, when the original videos are cropped, there is no overall difference in perceived contrast (some subjects even saw the magnified video as having higher contrast at this distance). When the original videos were not cropped, subjects were still reducing their contrast by about 7% to match the magnified videos.

4.5 Discussion

The decline with distance of the magnification effect on perceived contrast co-occurs with the decrease in visibility of the video pixelation, i.e. a decrease in perceived blur. There is a residual effect of magnification which remains even when the pixelation of the magnified video is invisible, but this effect is only seen when the comparison (original) video is at full-size. This is most likely caused by the mismatch in physical contrast between whole videos and their central subregions for the full size original condition, but the effect is still far smaller than the prediction made in Section 3.

5. General Discussion

We have confirmed that there is a decrease in the perceived contrast of video clips as a result of digital magnification. The magnitude of perceived attenuation is around 10-20% (Experiments 1 and 2) over the range of magnifications we used. This illusory attenuation of perceived contrast for magnified video clips can be attributed principally to the blurring or pixelation caused by magnification, as we found that it is largely removed when the pixel separation is made invisible (Experiment 3). We also find that the structure of the video background can have a significant effect on its perceived contrast (Experiment 2).

The perceived contrast of a complex visual stimulus amounts to pooling of contrasts across both spatiotemporal and frequency domains (Cannon & Fullenkamp 1991²⁰, Haun & Peli 2011¹⁷, Bex & Makous 2002⁸). *Spatiotemporal pooling* is implied by the effect of including content in the original resolution videos that is not magnified in the comparison videos – in every condition using this stimulus arrangement, the magnified videos were judged as being of lower contrast than the originals. This result is what would be expected if subjects were judging video contrast by pooling brightness and darkness estimates over the entire area of the stimulus videos. However, magnified video contrast was not underestimated relative to the originals to the degree predicted if subjects were directly comparing the global RMS contrasts of the two videos (i.e. data in Figure 5b did not track with the differences plotted in Figure 2) – this is not too surprising, since we should not expect that the RMS measure of contrast should directly determine perceived contrast. Neither is it surprising that subjects should pool their estimates of video contrast over the video area and duration, since this was what they were instructed to do. The structure of this

pooling is relatively unknown, although larger deviations from the local mean luminance – i.e. higher local contrasts – are likely to contribute inordinately to the pooled estimate.

Spatial frequency pooling is demonstrated by the existence of a perceived difference in contrast when the original and magnified videos contain identical content (the cropped conditions), and the disappearance or reversal of this difference at large viewing distances. This is explained by the visible lack of high-frequency contrasts in the magnified video, i.e. the perceived blurring – the effect disappears at large viewing distances because the lack of high-spatial-frequency contrast is no longer visible. If perceived contrast is a summation over a perceptually fixed range of spatial frequencies, a stimulus perceived as blurred will seem to have lower contrast even if its luminance structure is identical (apart from scale) to a sharp stimulus. The structure of this pooling is likely rather complex, since the precise relationship between perceived and physical contrast in a broadband image is dependent on spatial frequency (Haun & Peli 2012¹⁷). The most important feature of this dependency is the acuity limit: high enough spatial frequencies cannot be detected, and so cannot contribute to contrast judgments.

6. Summary

Magnified video is perceived as having lower contrast than normal-resolution video for two reasons: First, because regions of an image tend to have lower overall physical contrast than the larger image by conventional measures; second, and most importantly, because the magnified image appears blurred. This perceived (but not physical) blur is interpreted as loss of contrast as in many situations. When asked to compare video contrasts, subjects do not seem to be comparing the actual luminance distributions (RMS contrasts). Caution is to be taken in measuring perceived video contrast – and by that token perceived *quality* - against a background of fixed mean luminance, especially when video size is allowed to vary. Finally, the contrast attenuation we have identified may be specific to the basic interpolation method used. More sophisticated algorithms than bilinear interpolation are designed to preserve edge sharpness in the magnified image, and these would presumably also preserve image contrast. The method we have demonstrated in this paper can be used to compare the basic perceptual impact – the perceived contrast - of other interpolation algorithms.

Acknowledgments

Supported in part by NIH grants EY05957(EP), EY12890(EP) and EY19100(RW) and an unrestricted grant from Analog Devices Inc.

8. References

1. Farsiu S, Robinson D, Elad M, Milanfar P. Advances and challenges in super-resolution. *Int. J. Imaging Syst. Technol.* 2004; 14:47–57.
2. Goldstein RB, Apfelbaum H, Luo G, Peli E. Dynamic magnification of video for people with visual impairment. *SID Symposium Digest.* 2003; 37:1152–1155.
3. Peli E, Luo G, Bowers A, Rensing N. Development and evaluation of vision multiplexing devices for vision impairments. *Int. J. Artif. Intell. Tools.* 2009; 18:365–378. [PubMed: 20161449]
4. Knoche, H.; Papaleo, M.; Sasse, MA.; Vanelli-Corali, A. *ACM Multimedia Proceedings*. Vol. 15. Augsburg, Germany: 2007. The kindest cut: enhancing the user experience of mobile TV through adequate zooming; p. 87-96.
5. Peli E. In search of a contrast metric: Matching the perceived contrast of Gabor patches at different phases and bandwidths. *Vision Res.* 1997; 37:3217–3224. [PubMed: 9425539]
6. Peli E, Yang J, Goldstein R, Reeves A. Effect of luminance on suprathreshold contrast perception. *J. Opt. Soc. Am. A-Opt. Image Sci. Vis.* 1991; 8:1352–1359.
7. Wang, Z.; Bovik, AC. *Modern Image Quality Assessment*. Morgan & Claypool: 2006.

8. Bex PJ, Makous W. Spatial frequency, phase, and the contrast of natural images. *J. Opt. Soc. Am. A-Opt. Image Sci. Vis.* 2002; 19:1096–1106. [PubMed: 12049346]
9. May KA, Georgeson MA. Blurred edges look faint, and faint edges look sharp: The effect of a gradient threshold in a multi-scale edge coding model. *Vision Res.* 2007; 47:1705–1720. [PubMed: 17442367]
10. Watt RJ, Morgan MJ. The recognition and representation of edge blur - evidence for spatial primitives in human-vision. *Vision Res.* 1983; 23:1465–1477. [PubMed: 6666047]
11. Roufs, JAJ. Human Vision, Visual Processing, and Digital Display 66-72. SPIE - Int Soc Optical Engineering; Bellingham: 1989. Brightness contrast and sharpness, interactive factors in perceptual image quality.
12. Peli E. Contrast sensitivity function and image discrimination. *J. Opt. Soc. Am. A-Opt. Image Sci. Vis.* 2001; 18:283–293. [PubMed: 11205973]
13. Bex PJ, Dakin SC, Mareschal I. Critical band masking in optic flow. *Netw.-Comput. Neural Syst.* 2005; 16:261–284.
14. Brainard DH. The psychophysics toolbox. *Spatial Vis.* 1997; 10:433–436.
15. Pelli DG. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vis.* 1997; 10:437–442.
16. Loftus GR, Masson MEJ. Using confidence-intervals in within-subject designs. *Psychon. Bull. Rev.* 1994; 1:476–490.
17. Haun AM, Peli E. Measuring the perceived contrast of natural images. *SID Symposium Digest.* 42:302–304.
18. Chubb C, Sperling G, Solomon JA. Texture interactions determine perceived contrast. *Proc. Natl. Acad. Sci. U. S. A.* 1989; 86:9631–9635. [PubMed: 2594791]
19. Heinrich SP, Bach M. Less is more: Subjective detailedness depends on stimulus size. *J. Vision.* 2010; 10:10.
20. Cannon MW, Fullenkamp SC. A Transducer model for contrast perception. *Vision Res.* 1991; 31:983–998. [PubMed: 1858328]

Figure 1.

Stimulus configuration. On one side of the display, the magnified video was presented. On the other side, the original video was presented either in its entirety (spanning 360 pixels) or cropped to match the content of the magnified video. Videos looped continuously, forward and backward, until the subject indicated which video had higher contrast. The background/surround structure (not shown) was varied depending on experiment condition.

Figure 2.

RMS contrast of cropped/magnified video clips is plotted relative to the original video contrast. As magnification increases, contrast declines (circle symbols). Scrambling the phase spectrum of the original image lessens the effect of magnification (square symbols); whitening the amplitude spectrum without changing the phase spectrum decreases the effect even more (triangle symbols). White noise contrast does not decrease with magnification (X symbols). The dashed line is $1/m$, a relationship that closely describes the effect of cropping on video contrast.

Figure 3.

Luminance histograms for a cropped frame at its original resolution of 90x90 px (solid dots), and the same frame interpolated to four times its original size (X symbols). RMS contrast and mean luminance are indicated by the horizontal bars and their center points. Neither property is affected by the interpolation.

Figure 4.

Experiment 1: Effect of electronic magnification (abscissa) on perceived contrast (ordinate). Mean of 6 observers shown. Matching contrast is presented as a proportion of the magnified video contrast (which was clamped at 63% of its original value), the 'contrast estimate'. Paradoxically, the effect of magnification was smaller when the matching video included the whole image (full-size). The dashed line is the reciprocal of the square root of magnification factor, the expected average difference in RMS contrast between a magnified and full-size original video (as in Figure 2). Error bars are 95% confidence limits (Loftus & Masson 1994¹⁶).

Figure 5.

Matching ratios for all data from Experiments 1 and 2. Legend identifies conditions where the background luminance was fixed or varied; and where the background was blank or filled with Gaussian noise. Data are jittered slightly along the abscissa to make different conditions discriminable. a) When the originals were cropped to match the magnified content, there was a large effect of controlling the background structure, reducing the perceived difference in contrast between magnified and original videos. b) When the originals were unchanged across magnification levels, the background structure was less important, but there was still some effect. The dashed line is the RMS-difference prediction as shown in Figure 4. c) The difference between cropped and full-size conditions reveals the effect of cropping on the perceived contrast of the *original* videos. Except when the background is filled with noise and luminance-matched with the video, the effect of cropping is generally to increase the perceived contrast of the unmagnified video. Thus the effect sizes shown in a), except for the smallest effects (vary/noise), are exaggerated by cropping the original video. Error bars in a) and b) are 95% confidence limits (Loftus & Masson 1994¹⁶). Error bars in c) are the Pythagorean sum of the error bars of a) and b).

Figure 6.

a) Video pixel separation for original-resolution and magnified videos at the three viewing distances. Normal human vision limits resolution of detail to around 1 arcmin. b) Effect of magnification on perceived contrast of video at three viewing distances, for cropped and full-sized original comparison conditions. Magnification was 3x at all distances. Error bars are 95% confidence limits (Loftus & Masson 1994¹⁶).