PREDICTION OF PREFERRED CLEARTYPE FILTERS USING THE S-CIELAB METRIC

Jiajing Xu[1], *Joyce Farrell*[1], *Tanya Matskewich*[2], *and Brian Wandell*[1]

[1] Department of Electrical Engineering, Stanford University
[2] Advanced Reading Technology, Microsoft

ABSTRACT

The appearance of rendered text is a compromise between the designer's intent and the display capabilities. The ClearType rendering method is designed to enhance rendered text by exploiting the subpixel resolution available on color displays. ClearType represents the high-resolution font outline at the full subpixel resolution of the display and then filters the image to enhance contrast and reduce color artifacts. The filter choice influences text appearance, and people have clear preferences between the renderings with different filters. In this paper, we predict these preferences using S-CIELAB, a spatial extension to the perceptual color metric CIELAB. We calculate the S-CIELAB difference between designed and rendered fonts for various filters. We compare the size of these differences with preference data obtained from individual subjects.

***Index Terms*—** ClearType, S-CIELAB

1. INTRODUCTION

In most color displays, each pixel is composed of three horizontally adjacent subpixels that emit the red, green, and blue (RGB) primary lights. Traditional display algorithms treat the subpixels as spatially coincident and forfeit the potential resolution enhancement in the horizontal dimension. ClearType uses the individual subpixel elements in every pixel to increase the details in text display.

The attempt to increase horizontal resolution, however, may result in color fringing errors. Platt [1] quantified the color errors using a model inspired by psychophysical experiments and developed direct linear optimal filtering to minimize such error. This is done by applying three color filters to each color channel of the full-color input image to produce values for each subpixel, i.e. the three filters applied to red color channel are denoted by R→R, G→R, and B→R. Subsequently, Betrisey [2] found that the filters between different color channels (R→G, R→B, G→B, etc.) have relatively little power. The remaining three filters (R→R, G→G, and B→B) are nearly identical but centered at different subpixel. Hence, Betrisey replaced the nine filters with one box filter and referred this approximation as RGB decimation with displaced box filters. These simplifications are the basis of a real-time implementation of ClearType.

The rendered text is designed to appeal to the human viewer. The visibility of the color artifacts and contrast of the text will depend on the display characteristics, the filtering method, and properties of the human visual system. In this paper we combine device simulation and perceptual metrics (S-CIELAB) to predict the visibility of unwanted artifacts in ClearType text. We test the idea that these artifacts are at the heart of users' preferences by comparing the artifact visibility with perceptual preference data.

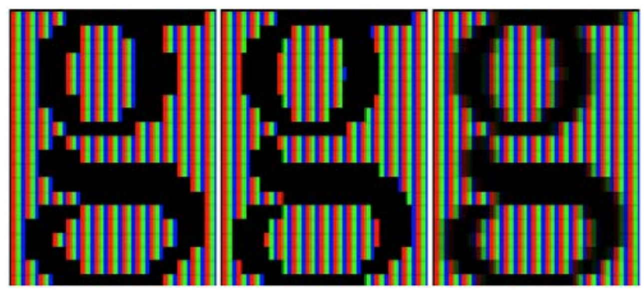

(A) (B) (C)

Figure 1: (A) True-type character rendered at full pixel resolution. (B) The same character rendered at subpixel resolution. (C) The subpixel rendering filtered horizontally. Simulated font: 10-pt 'g' from Georgia on a 200-dpi display. The DST estimates a full radiance image (not shown).

2. GENERATING CLEARTYPE FONTS

A rasterizer converts text from a scalable, often continuous, font description (such as TrueType fonts) to a bitmap format that specifies each pixel as on or off (Figure 1a). Many rasterizers simply specify whether all the subpixels should be on or off. The ClearType rasterizer converts the font description to the resolution of the colored subpixel grid, which has three times the horizontal pixel resolution (Figure 1b). Displayed without further modification, the rendered text would have higher horizontal resolution but the subpixel colors would be visible. To reduce the visibility of the unwanted color fringes, the raster image is filtered in the

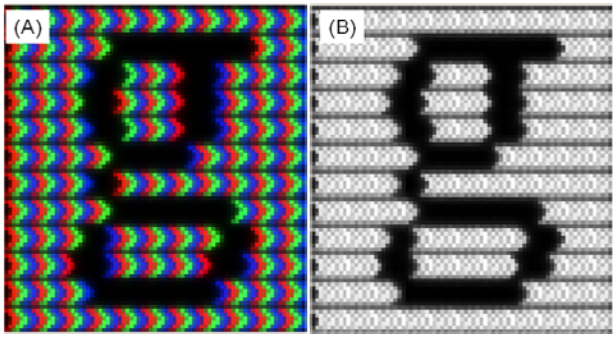

Figure 2: The raster images of ClearType fonts without filtering (A) on a color display, and (B) on a matched-monochrome display. The calibrated display is a 98-dpi Dell LCD Model 1905P, which has chevron-shaped RGB subpixel structure.

horizontal direction (Figure 1c). The analyses in this paper create tools to simulate the radiance image of the rendered text and to predict user preference for different ClearType filters.

3. DISPLAY SIMULATION

To predict the appearance of different images, we require a physical description of the radiance image. We use the Display Simulation Toolbox (DST) to create the radiance image of a ClearType font. The DST uses the measured subpixel point spread functions, spectral power distribution, and gamma response functions to calculate the display radiance. Elsewhere, we have shown that for certain displays the prediction is as accurate as the magnified images of the characters captured by a calibrated high resolution imaging system. [4] Figure 2 shows the predicted full luminance image estimated for a calibrated LCD panel that consists of RGB subpixels in a chevron configuration. The data are simulated at matched spatial resolutions.

3. PREDICTING PREFERENCE

We predict user preference by comparing the ClearType rendering with text rendered on a matched display, in which all of the subpixels are white. In this case, the font outline on the matched display is identical to the outline used in the ClearType rendering and the differences are all due to contrast and color. We compute the visibility of the difference using S-CIELAB, a spatial extension of the CIELAB metric [3]. The CIELAB color metric was designed for large uniform color fields. S-CIELAB added a spatial processing step prior to the CIELAB calculation to simulate the spatial blurring of human vision system. Since the fonts are patterned images, S-CIELAB is the appropriate metric. We use the mean S-CIELAB ΔE across the whole image to measure the visibility of the difference between the ClearType rendering and the rendering on a matched monochrome display and we assume that users prefer the rendering that is closest to the image on that display.

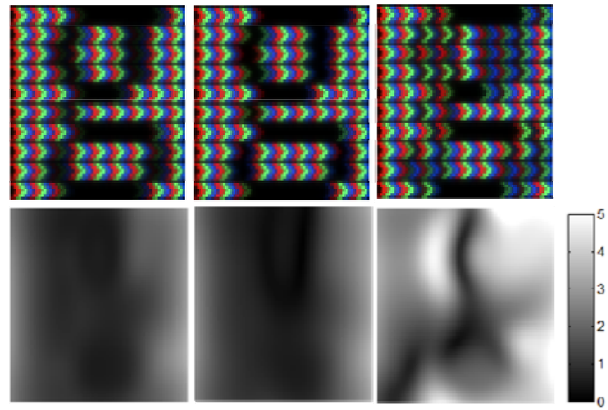

Figure 3: Comparison of three filter kernels. First row: filtered text; Second row: S-CIELAB ΔE map. Filter parameters (*a, b*) from left to right are (0.1, 0.2), (0, 0.3) and (0.4, 0.4).

4. RESULTS

First, we calibrated a display to determine the spectral radiance of its primaries, the spatial spread of the subpixels and to determine whether the screen radiance could be predicted by a model including only static nonlinearities [4]. These conditions were satisfied by the Dell LCD panel model 1905P. Then, we measured the S-CIELAB difference between the simulated radiance images of text rendered on the calibrated display and text rendered on a display which is matched to the calibrated display but containing only white subpixels.

We measured the difference for a series of ClearType filters. Specifically, we evaluated five-tap ClearType filters, [a b c b a], where $2a+2b+c=1$. The filters are parameterized by the (a, b) values, and each combination results in different color fringing errors for the ClearType fonts. Figure 3 shows a ClearType font glyph filtered with three different filter kernels. The S-CIELAB difference (ΔE) between the simulated radiance images of text rendered on the calibrated display and text rendered on an idealized display is shown as an error (ΔE) map. As the (a, b) values change, the ΔE map and mean ΔE value change significantly. These errors capture changes in the rendered font colorfulness and contrast.

We performed a grid search across a range of (a, b) values to find the filter kernel that minimizes the difference between the font rendered on the calibrated display and on the matched ideal display. Figure 4 illustrates the analysis for the lowercase letter 'g', size 10, from Georgia font family (g-Georgia-10 in short). The simulated viewing distance is 0.38m (15 inches), a typical distance. The calibrated display is a Dell LCD monitor that has chevron-shaped RGB subpixels (98 dpi). Figure 4a shows a contour plot of the mean

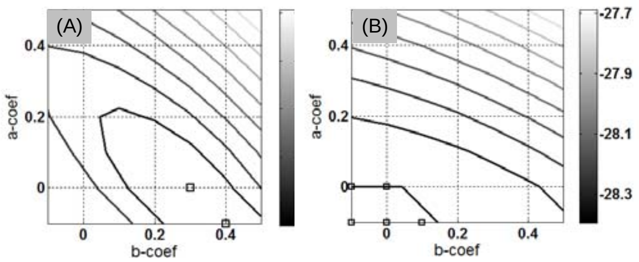

Figure 4: Contour plots showing S-CIELAB (A) and PSNR (B) errors as a function of filter coefficients, (*a*, *b*). Squares show values with minimum error. Preference contours simulated for g-Georgia-10 on a Dell 1905P.

S-CIELAB ΔE over a range of filter parameters (*a*, *b*). The global minimum of ΔE=1.25 was found near *a*=0 and *b*=0.3, which is a three-tap filter. A range of filter kernels a comparable ΔE value.

We performed the same analysis using the PSNR metric. The contour maps for this metric are shown in Figure 4b. The smallest error occurs with no filtering at the parameter values (*a*=*b*=0).

The PSNR analysis has several undesirable features [5]. First, the PSNR predictions are independent of viewing distance and display properties. Second, the PSNR metric does not predict preference well in this application, nor does it match the preference measurements described below. We include the PSNR analysis for comparison with S-CIELAB because PSNR is used widely in the engineering literature.

5. COMPARISON OF S-CIELAB PREDICTIONS AND SUBJECTIVE PREFERENCE DATA

In a related study, Farrell et al. [6] performed a series of visual psychophysical experiments to determine the preferred ClearType filter. In the experiment, different versions of the ClearType characters were created by varying the two filter parameters *a* and *b*. Subjects were asked to indicate which version, compared with a series along a line in (*a*, *b*) space, they preferred. Subjects made this choice for many possible lines. In this section, we compare the results obtained from these preference experiments with predictions based on the S-CIELAB predictions of artifact visibility.

Figure 5 shows the predictions for g-Georgia-10 with different ClearType filtering presented on a Dell LCD with chevron-shaped subpixel [4]. The iso-preference contours are shown by the dashed-line contour plots. Each contour falls off with the number of times the subject indicated a preference for that filtered version. The contour plots for the preference data and the S-CIELAB predictions are similar, although not exact. The contour plot for the PSNR predictions (Figure 4b) differ significantly from the preference data. In particular, the PSNR prediction, (0, 0) is never preferred by the subjects.

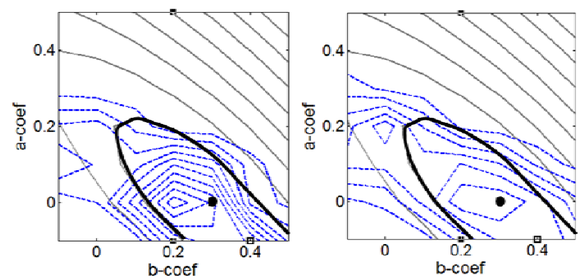

Figure 5: Comparison of the predicted iso-preference contours using S-CIELAB (solid gray lines) and measured iso-preference contours [6] (dashed blue lines). The two panels are data from two subjects. The darkened contour is the S-CIELAB predicted peak preference contour. The filled circle is the estimated preferred filter. Simulated for g-Georgia-10 on a Dell 1905P.

The S-CIELAB predictions and the preference data both suggest that a range of filter parameters (*a*, *b)* yield equivalent preferences. This range includes values when *a*=0, and thus the five-tap filter becomes a three-tap filter.

The S-CIELAB predictions vary with display model and viewing distance. Under typical viewing conditions, the combination of *a*=0 and *b*=0.3 is often close to the preferred filter setting. However, these numbers can vary significantly between displays and viewing conditions.

In generating the filter preference maps for this display and font, we always find a number of filters which result in very low color fringing errors. This suggests that we might find a ClearType filter that works reasonably well at reducing the visibility of color fringing errors for all characters. We calculated the filter performance map for several 10-point English letters from Georgia font family rendered on the Dell 1905P. The predicted differences were near (0, 0.3).

The predicted preference shifts when we perform simulations using other displays or viewing conditions. At a larger viewing distance (e.g. 1.5m), the predicted preference is near (0, 0). In the cases we have explored, (e.g., LCD stripe), the preferred region is generally defined by a zone in which (*a*, *b*) are negatively correlated, falling near a line such as *a*+*b*=0.6. We are testing these predictions with additional user preference data.

6. DISCUSSION

The metric we use is a “full-reference image quality metric” that compares a test image with an ideal reference image. In this case, the test image is a ClearType font character rendered on a calibrated color display and the ideal reference image is the same character rendered on a monochrome display with spatial resolution matched to the subpixel resolution of the color display. Because the font outline on the reference monochrome display is identical to the font outline on the color display, the differences between the reference and test images are all due to contrast and color. Our analysis differs from previously published full-reference image quality metrics [7,8] in several important ways. First, we calculate the full radiance of both the reference and test images. Second, we calculate the perceptual spatial-chromatic difference between the two radiance images. Third, since our stimuli are isolated characters, we do not consider the effects of spatial masking.

A promising new approach to image quality assessment is the development of no-reference image quality metrics that quantify the visibility of annoying artifacts, such as blur [9] and JPEG blocking artifacts [10,11]. We are currently exploring the application of a reference-free metric to predict the visibility of color artifacts introduced by non-preferred ClearType filter parameters.

7. SUMMARY

We describe a collection of methods that can be used to predict preferences among different ClearType filters. These methods begin with a specification of the radiance image of the display. We further assume that users have an implicit ideal font, which is captured by a ClearType rendering on a monochrome display with spatial resolution matched to the subpixel resolution of the test display. We find that the S-CIELAB differences between the test display and this idealized monochrome display predicts the pattern of choices of user preference reasonably well, and significantly better than the PSNR metric. A range of filter kernels reduce the visibility of color fringing, and the range of these parameters corresponds fairly well to user preferences for the same display and viewing conditions. The methodology is open to further empirical testing and comparison with user preference data. Through these measurements, we should be able to evaluate how well the radiance image and S-CIELAB metric predicts user preferences for ClearType filters.

8. ACKNOWLEDGEMENTS

We thank Greg Hitchcock and Kevin Larson for their interest and support for this research.

9. REFERENCES

[1] Platt, J., "Optimal Filtering for Patterned Displays," *IEEE Signal Processing Letters*, 7, 7, pp. 179-180, (2000).

[2] Betrisey, C., Blinn, J. F., Dresevic, B., Hill, B., Hitchcock, G., Keely, B., Mitchell, D. P., Platt, J. C., Whitted, T., "Displaced Filtering for Patterned Displays," *Proc. Society for Information Display Symposium*, pp. 296-299, (2000).

[3] X. M. Zhang and B. A. Wandell, "A spatial extension to CIELAB for digital color image reproduction," *Proc. Society for Information Display Symposium*, San Diego, CA., 1996.

[4] Farrell, J., Ng, G., Ding, X., Larson, K. and Wandell, B., "A Display Simulation Toolbox for Image Quality Evaluation," *IEEE/OSA Journal of Display Technology*, vol. 4, no. 2, pp. 262-270, June 2008.

[5] B. Girod, "What's wrong with mean-squared error," in *Digital Images and Human Vision*, A. B. Watson, Ed. Cambridge, MA: MIT Press, 1993, pp. 207–220.

[6] Farrell, J. E., Larson, K., Matskewich, T., and Wandell, B., "Subjective preference for ClearType Characters," manuscript in preparation.

[7] S. Daly, "The visible difference predictor: an algorithm for the assessment of image fidelity," in *Digital Images and Human Vision*, A. B. Watson, Ed. Cambridge, MA: MIT Press, 1993, pp. 179–206.

[8] J. Lubin, "The use of psychophysical data and models in the analysis of display system performance," in *Digital Images and Human Vision*, A. B. Watson, Ed. Cambridge, MA: MIT Press, 1993, pp. 163–178.

[9] P. Marziliano, F. Dufaux, S. Winkler, and T.Ebrahimi, "A No-Reference Perceptual Blur Metric," *IEEE Int. Conf. on Image Processing*, pp. 57-60, Sep. 2002.

[10] Z. Wang, A. C. Bovik, and B. L. Evans, "Blind Measurement of Blocking Artifacts in Images," *IEEE Int. Conf. on Image Processing*, pp. 981-984, Sept., 2000.

[11] Z. Wang, H. R. Sheikh, and A. C. Bovik, "No-Reference Perceptual Quality Assessment of JPEG Compressed Images," *IEEE Int. Conf. on Image Processing*, pp. 477-480, Sept. 2002.