Christopher Tsai
April 15, 2007
EE 368 – Digital Image Processing

Problem Set I

Problem 1 – Quantization

First, let us concentrate on the illustrious Lena:

Problem 1A - Original Lena Image

Problem 1A - Quantized Lena Image

Problem 1B - Dithered Lena Image

Problem 1B - Dithered and Quantized Lena Image

Problem 1D - Lowpass Filtered Quantized Lena Image

Problem 1D - Lowpass Dithered & Quantized Lena

IMAGE PROCESSING PROCEDURE	LENA
Original Image	0
Eight-Level Quantized Image	83.760159
Dithered Image	85.485791
Dithered & Quantized Image	167.337551
Lowpass-Filtered Quantized Image	62.355633
Lowpass-Filtered Dithered & Quantized Image	48.730312

Because our image contains 256 levels by default, quantizing the image into eight evenly spaced intervals restricts the number of possible pixel values to eight: 16, 48, 80, 112, 144, 176, 208, and 240. This quantization introduces not only a finite mean squared error (83.76) but also sets of contours in the image connecting pixels that seemed identical in the original image; even though these contoured regions originally belonged to a larger set of seemingly uniform gray values, quantization pushed them far enough apart that they formed new regions, which deeply disturb our perception of the image. These additional areas protrude noticeably from Lena's hat and skin, giving the image an artificial, contrived appearance.

In an effort to eliminate these contours, we dither the image by adding random uniformly distributed noise, with half-level amplitudes sufficient enough to push pixels out of contoured regions during quantization. In other words, with small amounts of random noise added to break up the total uniformity of pixel values in any given region, we randomize quantization error and reduce contouring to such a point-wise effect that the groups never form large enough to displease our eyes. However, even without quantization, we see that the noise added during dithering introduces mean squared error (MSE $\approx$ 85.486) nearly equivalent to quantization itself. The noise manifests itself as small grains speckled haphazardly throughout the image. Quantizing the dithered image increases the mean squared error even further (167.338), and the graininess resulting from noise becomes even more visible. This degeneration occurs because the noise distances the processed image from the original image and contributes to the mean squared error at every point; subsequent quantization further distances points from their original values since the noise could have changed pixel values beyond the threshold between quantization levels, thus increasing the mean squared error two-fold.

Clearly, the dithered and quantized image yields a much higher mean-squared error than the merely quantized image, so Part (B.) would appear to be a more erroneous image. However, because dithering effectively eliminates visible contouring from the quantized image, the result in Part (B.), despite its higher mean squared error, is more visually pleasing than the quantization in Part (A.), whose contours degrade the image's subjective quality to the sensitive human eye.

While Part (B.) wins only in visual quality after quantization, lowpass-filtering the results of Part (A.) and Part (B.) lead to a clear-cut victory for Part (B.) not only in visual quality but also in mean squared error; lowpass-filtering blurs the image and smoothens edges but fails to remove contouring, so Part (A.) still looks worse than Part (B.). Furthermore, the smoothening tempers the effect of the dithering noise by smoothening over the sharp fluctuations, thus improving the visual

quality of the dithered and quantized image from Part (B.). Most importantly, however, dithering significantly decreases the mean squared error by blending the noise into the image. Now, pure quantization followed by filtering yields an MSE of 62.356, whereas the lowpass filter reduces the MSE of the dithered and filtered image to approximately 48.73, lower than any of the other mean squared errors. Now, not only the subjective quality but also the mean squared error of the dithered result in Part (B.) emerge superior to the quality and accuracy of the non-dithered image. Smoothening the noise greatly lowers its effect on the MSE, to the point that it beats even the filtered noiseless image of Part (A.)!

We now proceed to the Einstein image:

Problem 1A - Original Einstein Image

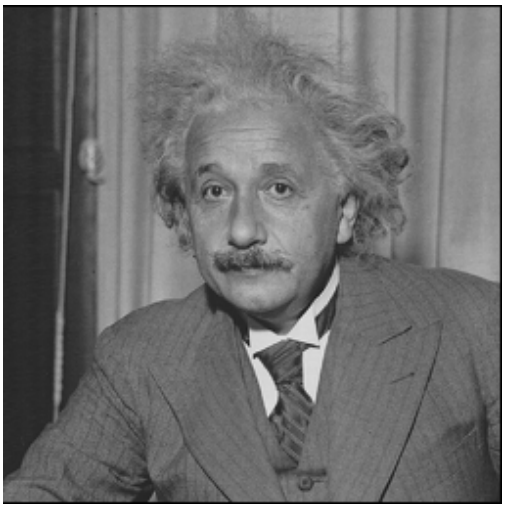

Problem 1A - Quantized Einstein Image

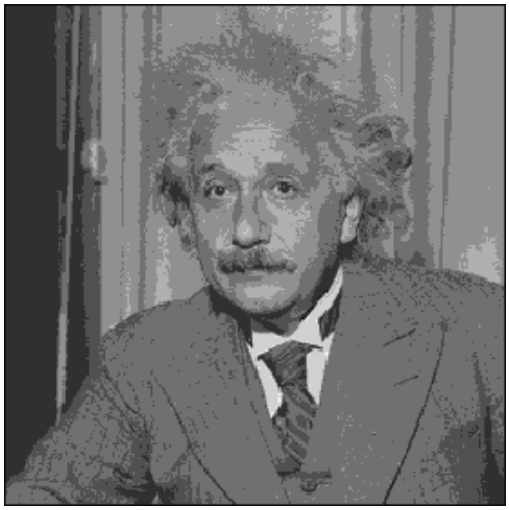

Problem 1B - Dithered Einstein Image

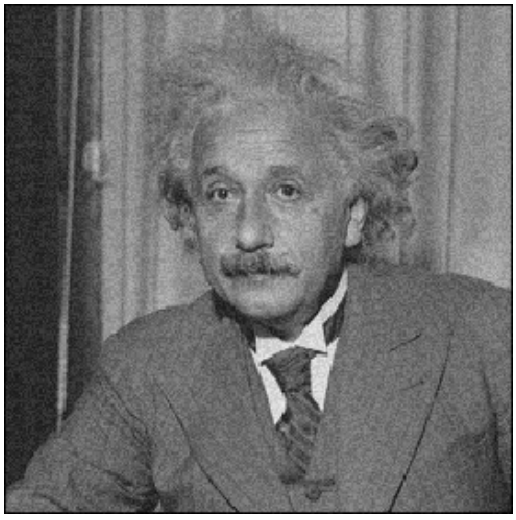

Problem 1B - Dithered & Quantized Einstein

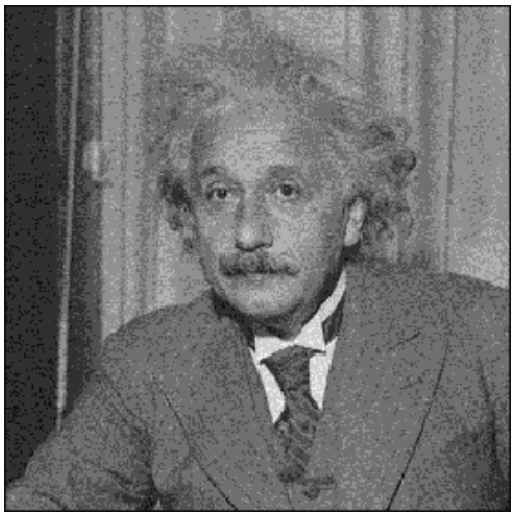

Problem 1D - Lowpass Quantized Einstein

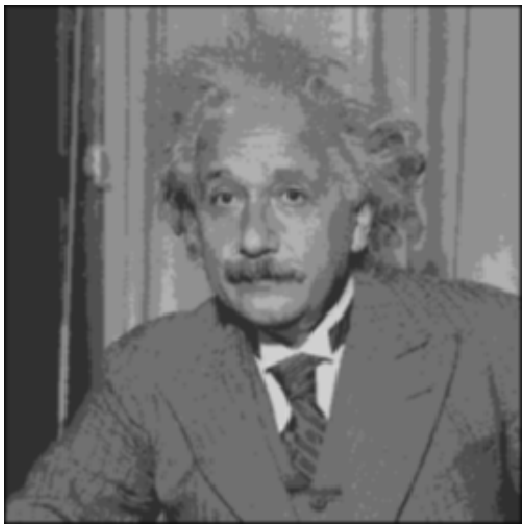

Problem 1D - Lowpass Dithered Quantized Einstein

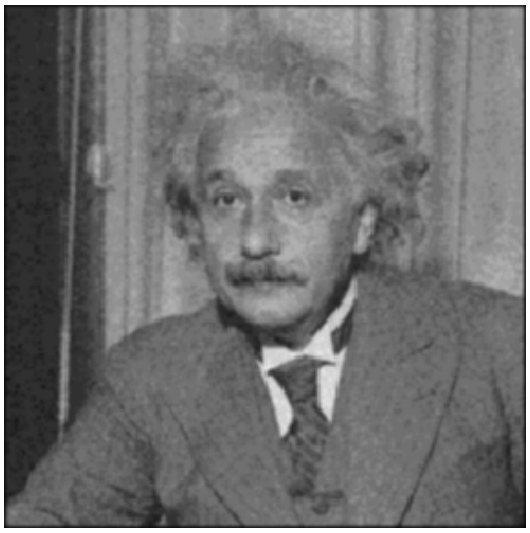

IMAGE PROCESSING PROCEDURE	LENA	EINSTEIN
Original Image	0	0
Eight-Level Quantized Image	83.760159	71.597488
Dithered Image	85.485791	84.506219
Dithered & Quantized Image	167.337551	153.997879
Lowpass-Filtered Quantized Image	62.355633	78.973221
Lowpass-Filtered Dithered & Quantized Image	48.730312	77.005875

As we observed in Lena, image quantization introduces undesirable contours, this time blending Einstein's forehead and hair into an unnatural amalgam of gray. Meanwhile, dithering the Einstein image prior to quantization yields a visible graininess throughout the image, but the random additive noise breaks up the contours, making the quantized image more visually pleasing, even if the mean squared error (153.998) is more than twice the mean squared error from Part (A.)'s noiseless quantization.

Unlike Lena, however, the dithered Einstein does not improve as significantly in MSE following lowpass-filtering. The lowpass filtering does improve the dithered image more than it improves the purely quantized image, as the contouring from Part (A.) remains, but the final MSE values of the noised and noiseless lowpass images are comparable (around 77-79). In fact, lowpass

filtering actually *increases* the MSE of the quantized image, perhaps because smoothening an image with so many sharp edges actually degrades accuracy; meanwhile, though, lowpass filtering following dithering and quantization continues to prove a powerful combination, as the MSE drops from its prodigious post-quantization value.

Finally, we conclude our image study with the Man:

Problem 1A - Original Man Image

Problem 1A - Quantized Man Image

Problem 1B - Dithered Man Image

Problem 1B - Dithered & Quantized Man Image

Problem 1D - Lowpass Filtered Quantized Man Image

Problem 1D - Lowpass Filtered Dithered & Quantized Man

IMAGE PROCESSING PROCEDURE	LENA	EINSTEIN	MAN
Original Image	0	0	0
Eight-Level Quantized Image	83.760159	71.597488	93.386036
Dithered Image	85.485791	84.506219	80.199482
Dithered & Quantized Image	167.337551	153.997879	164.101856
Lowpass-Filtered Quantized Image	62.355633	78.973221	129.470234
Lowpass-Filtered Dithered & Quantized Image	48.730312	77.005875	133.802326

Quantizing the Man image generates contours along the back and skin of the woman sitting behind the Man, but the visual displeasure resulting from contouring is not nearly as pronounced as it was for the Lena and Einstein images, which contained significantly more skin tones and background than the Man image.

As expected, dithering the image introduces speckle throughout, and quantization contouring vanishes as the random noise breaks up the uniformity of the woman's back and arm. Just as we concluded with Lena and Einstein, the subjective visual quality of the dithered quantized image dwarves the contoured Man image from Part (A.), even though the MSE of the dithered quantization is unfortunately much higher than the quantization MSE.

Lowpass-filtering the Man image brings the dithered quantization and the noiseless quantization to equal footing in terms of mean squared error, but neither MSE decreases below 120; the image error is still worse than the simple quantization case. The Man image responds less favorably to smoothing because the photo contains much more detail, from the rapidly changing colors in the background shirts and the detailed patterns on clothing to the straw in the lower-right hand corner. Thus, lowpass filtering (smoothening) blurs some of these details and therefore cannot achieve the low MSE we attained by filtering Lena and Einstein. However, when comparing the two lowpass-filtered images, we still prefer the dithered image since it contains no contours.

Problem 2 – Color Balancing

FRUIT IMAGE ILLUMINATION	GRAY-WORLD			SCALE-BY-MAX		
	k_R	k_G	k_B	k_R	k_G	k_B
Daylight	0.66905	1	2.823543	1.101184	1.143017	1.80742
Fluorescent Light	0.798162	1	5.52549	1.273973	1.083686	3.196875
Tungsten Light	0.433268	1	4.808248	1	1.540663	3.138037

Applying these algorithmic scale factors to their respective color slices, we obtain color-balanced images, which we display after compensating for CRT gamma nonlinearity with a $\frac{1}{\gamma}$ exponent:

Problem 2 - Original Fruit in Daylight Image

Problem 2A - Gray-World Fruit in Daylight

Problem 2B - Scale-by-Max Fruit in Daylight

Problem 2 - Original Fruit in Fluorescent Light

Problem 2A - Gray-World Fruit in Fluorescent

Problem 2B - Scale-by-Max Fruit in Fluorescent

Problem 2 - Original Fruit in Tungsten Light

Problem 2A - Gray-World Fruit in Tungsten

Problem 2B - Scale-by-Max Fruit in Tungsten

Even though each illumination source supplies a different spectrum of light with different perceived brightnesses, the two color-balancing algorithms yield similar effect on the fruit image. For example, consider the daylight-illuminated fruit with the following histogram distributions:

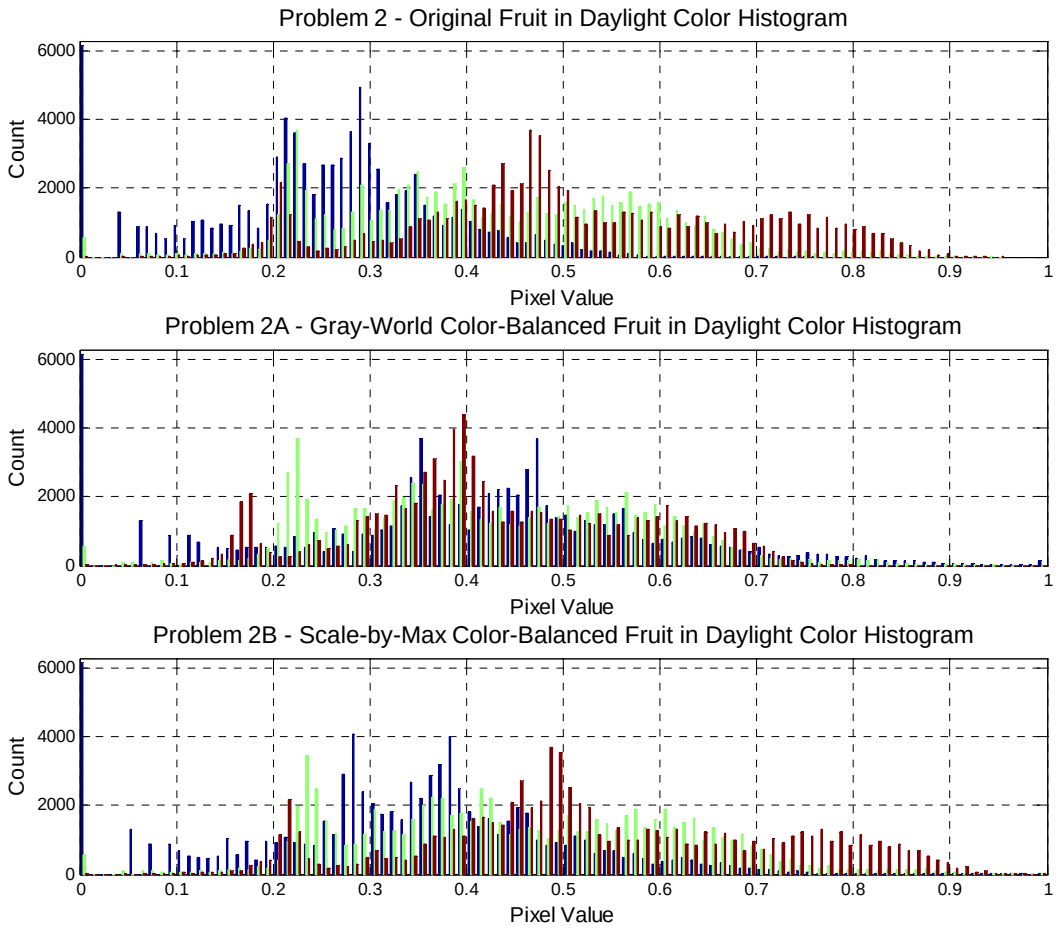

From the histogram, we see significant amounts of red with an entire range of intensities, most likely the result of the orange, pumpkin, red pepper, and apple. Green also appears liberally throughout the image, though less intensely than red. Finally, although blue appears to peak around 0.3, a significant number of pixels contain no blue component at all, as the high histogram bar at zero intensity reveals.

The Gray-World color-balancing algorithm yields a strangely colored image. As desired, we have preserved the mean of the green intensity values, but the other two colors now share coinciding density functions because of the gray-world assumption. As a result, the image of the fruit looks considerably greener and bluer than it should, with an unnatural aquamarine tint covering the entire image; this

displeasing shade arises from the disturbance of relative color proportions, as the red no longer has a higher mean intensity because of the gray-world normalization, which shifted all means to the green mean. Without relatively higher intensity in the red component and relatively lower intensity in the green component, the orange and red fruits and vegetables lose their natural expected colors.

The Scale-by-Max algorithm, on the other hand, preserves the peak locations of the three color primaries, so that the red component still possesses pixels with the highest intensity. The algorithm appears to have normalized the three distributions without merging their average values, so the resulting image color retains its relative proportions with intensities properly tuned. Both the histogram and the image ascertain that the Scale-by-Max algorithm performs best from our subjective perspective.

We now consider the fluorescent-lit fruit image, which features an intensity distribution with approximately equally centered blue and green components (≈0.2), with considerably brighter red pixels:

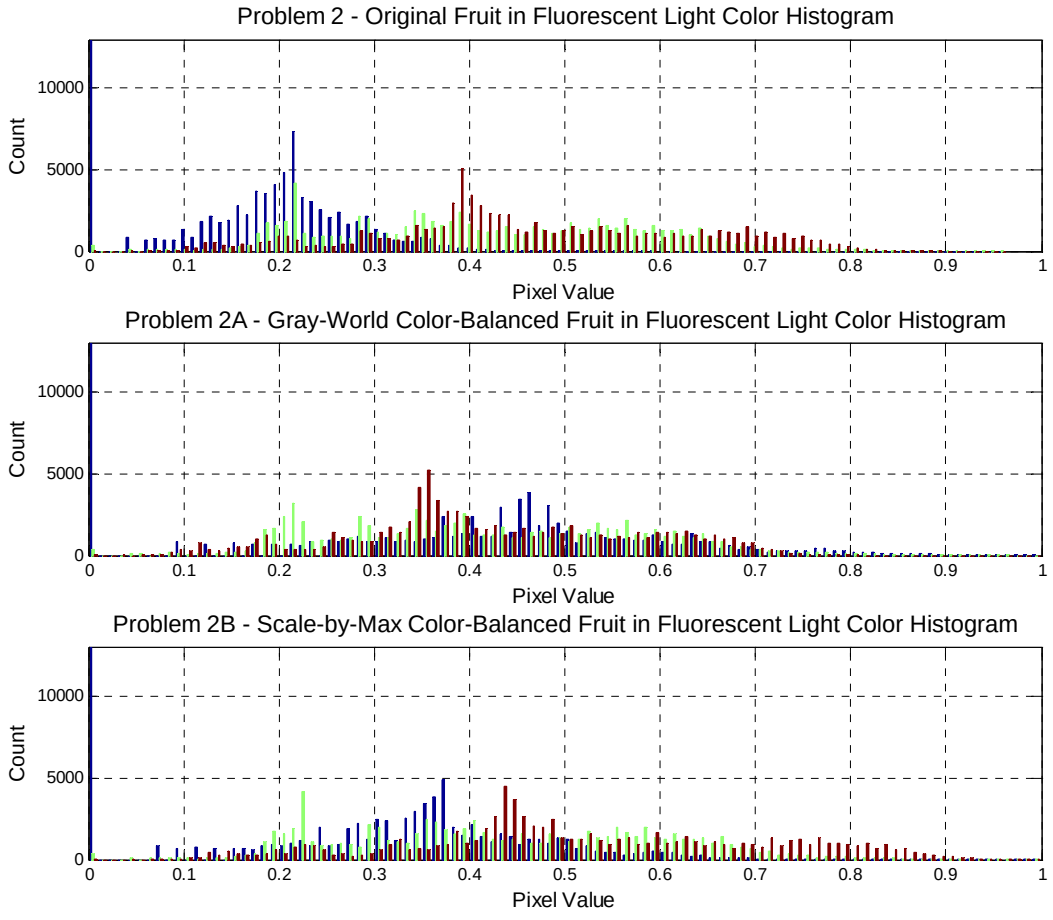

As with the daylight illumination, the Gray-World algorithm forces all three component distributions to behave with similar pattern and identical means, so the undesirable aqua-blue tint once again materializes. The histogram reveals that the red component, previously centered at much greater intensities, has been scaled below the other two components, explaining the seeming dearth of oranges and reds (or the seeming plethora of bluish green) in the Gray-World balanced image; hence, the Gray-World algorithm has once again disturbed the relative intensities of the three components, and, although the colors are balanced from a frequency standpoint, the relative intensities ultimately control our visual perception of the scene. The Scale-by-Max algorithm once again prevails in fluorescent light, with the frequencies scaled to match (except at zero intensity) and the relative intensities preserved.

The tungsten-lit fruit image responds similarly to the two algorithms:

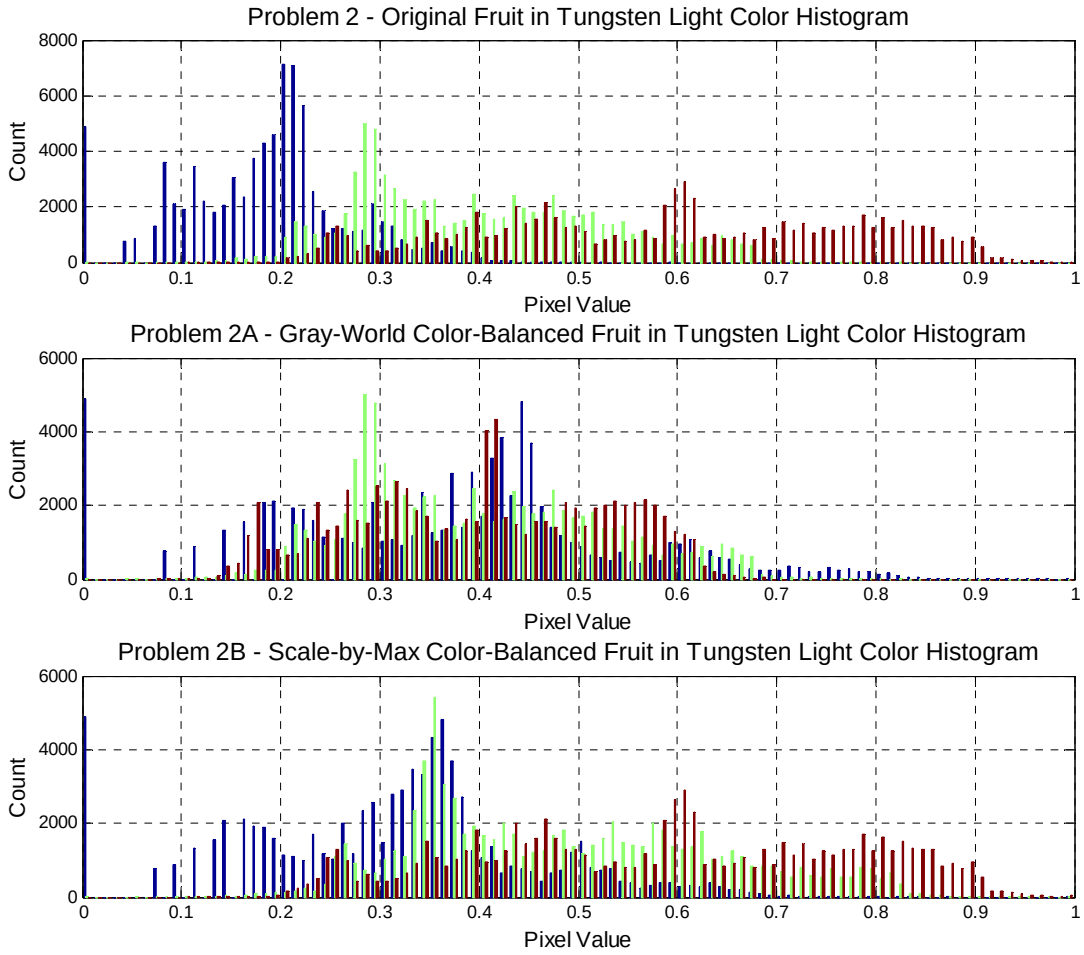

Unlike the other two images, the tungsten-lit scene contains considerably more blue pixels, which we can witness in the image as a greater brightness, or purer whiteness throughout the scene. Meanwhile, though red pixels are still the most intense, fewer of them cluster around a center main lobe. As we can tell from the histogram, the Gray-World algorithm succeeds in clustering the three distributions around the green mean, leading once again to the red imbalance and consequent bluish tint. Interestingly, the Scale-by-Max algorithm brings the green and blue central peaks to the same brightness (≈0.35), but the prevalence of intensity in the red component remains, preserving the natural balance of visible colors in the image.

Subjectively, our conclusion is no different; the Scale-by-Max algorithm performs better than the Gray-World algorithm because the resulting image looks the most genuine when juxtaposed next to the original unscaled image; the final tungsten image might look brighter than the other illumination final images, but the Scale-by-Max algorithm has approximately balanced the color frequencies, and the brighter whiteness results simply from a higher overall intensity distribution.

Because of its dependence on the mean value of a component, applying the Gray-World algorithm to the gamma pre-distorted images would not yield the same color-balanced images. Since gamma pre-distortion occurs as exponentiation by $\frac{1}{\gamma}$, the resultant color-balancing scale factors would become

$$k_R' = \frac{E\left[\{f_G(x,y)\}^{\frac{1}{\gamma}}\right]}{E\left[\{f_R(x,y)\}^{\frac{1}{\gamma}}\right]}$$

Each image point $[f_R(x,y)]^{\frac{1}{\gamma}}$ remains the same under pre-distortion and post-distortion, but its product with the new scale factor will not, in general ($\gamma \neq 1$), match our current post-distorted scale factor:

$$k_R = \left(\frac{\mu_G}{\mu_R}\right)^{\frac{1}{\gamma}} = \frac{\mu_G^{\frac{1}{\gamma}}}{\mu_R^{\frac{1}{\gamma}}} = \frac{\{E[f_G(x,y)]\}^{\frac{1}{\gamma}}}{\{E[f_R(x,y)]\}^{\frac{1}{\gamma}}} \neq \frac{E\left[\{f_G(x,y)\}^{\frac{1}{\gamma}}\right]}{E\left[\{f_R(x,y)\}^{\frac{1}{\gamma}}\right]}$$

$$k_R \neq k_R'$$

Therefore, the Gray-Scale algorithm will change depending on whether it operates before or after gamma distortion.

Conversely, as its name suggests, the Scale-by-Max algorithm computes its color-balancing scale factors by scaling by the *maximum* pixel value, thus making pre-distortion and post-distortion essentially identical. In other words, the new scale factor, computed with pre-distorted pixel values, will be

$$k_R' = \frac{1}{\max\left\{[f_R(x,y)]^{\frac{1}{\gamma}}\right\}} = \frac{1}{\left[\max\{f_R(x,y)\}\right]^{\frac{1}{\gamma}}} = \left[\frac{1}{\max\{f_R(x,y)\}}\right]^{\frac{1}{\gamma}} = {k_R}^{\frac{1}{\gamma}}$$

Because raising a number to a positive exponential is a monotonically increasing transformation, the maximum value of the pre-distorted pixel is equal to the post-distortion of the maximum undistorted pixel value. Finally, as with the Gray-Scale algorithm, distorting an image pixel before or after scaling yields no difference, since

$$k_R' \cdot [f_R(x,y)]^{\frac{1}{\gamma}} = {k_R}^{\frac{1}{\gamma}} \cdot [f_R(x,y)]^{\frac{1}{\gamma}} = [k_R \cdot f_R(x,y)]^{\frac{1}{\gamma}}$$

In conclusion, whereas the Gray-Scale algorithm generates different color-balanced images depending on the time of its application relative to gamma distortion, the Scale-by-Max algorithm will yield the same color-balanced product for both pre-distorted and post-distorted images; the Scale-by-Max procedure is gamma-distortion-invariant!