Christopher Tsai
February 2, 2008
EE 362 – Applied Vision and Image Systems

Tutorial III – Color Metrics

Problem #1 – Linearity in LAB Coordinates

First, let us consider the equation for L*, a*, and b* in LAB-coordinate space:

$$L^* = 116 \cdot \sqrt[3]{\frac{Y}{Y_W}} - 16$$

$$a^* = 500 \cdot \left(\sqrt[3]{\frac{X}{X_W}} - \sqrt[3]{\frac{Y}{Y_W}} \right)$$

$$b^* = 200 \cdot \left(\sqrt[3]{\frac{Y}{Y_W}} - \sqrt[3]{\frac{Z}{Z_W}} \right)$$

Because the LAB coordinates vary as the cube root of the luminance (and also X, Z) coordinates, equal spacing in LAB coordinates requires cubic spacing in reflectance. Thus, we select gray series reflectance levels with cubically increasing (or decreasing) spacing between their luminance (Y) values, such as the example curve plotted below:

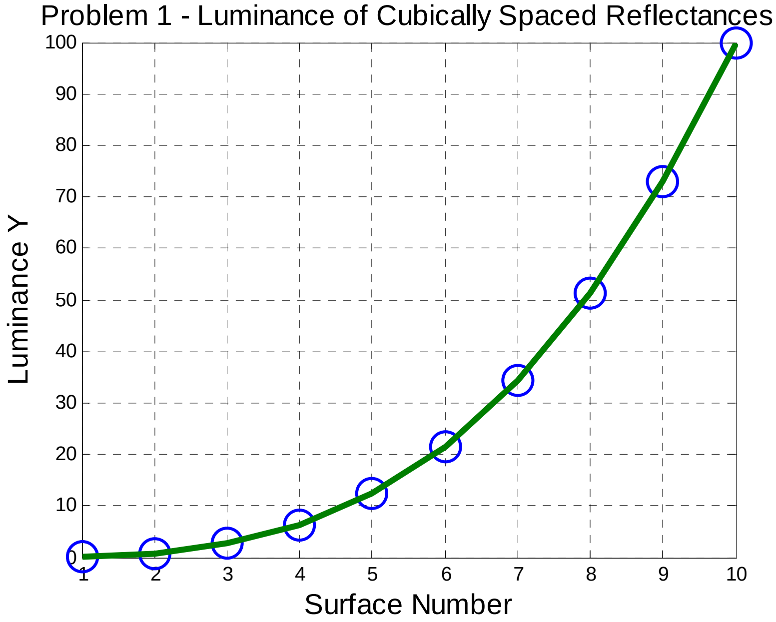

In the greater scheme of XYZ color space, the gray levels fall on a curve with very large cluster for low values of Y but much wider spacing between high-reflectance gray levels. Mathematically, we vary the spacing cubically instead of linearly:

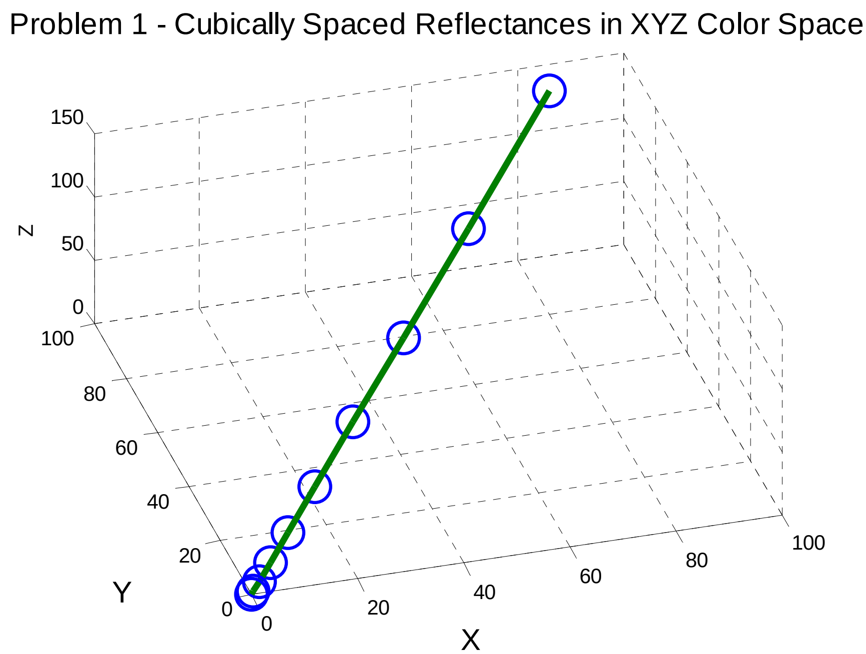

As a result of cubic spacing in reflectance, we obtain even (linear) spacing in LAB coordinates:

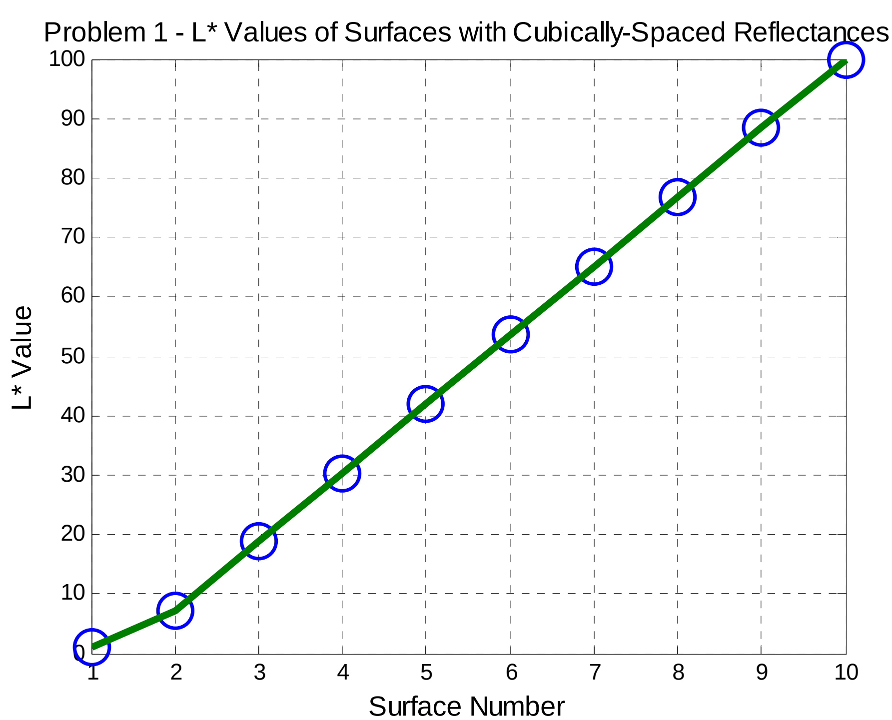

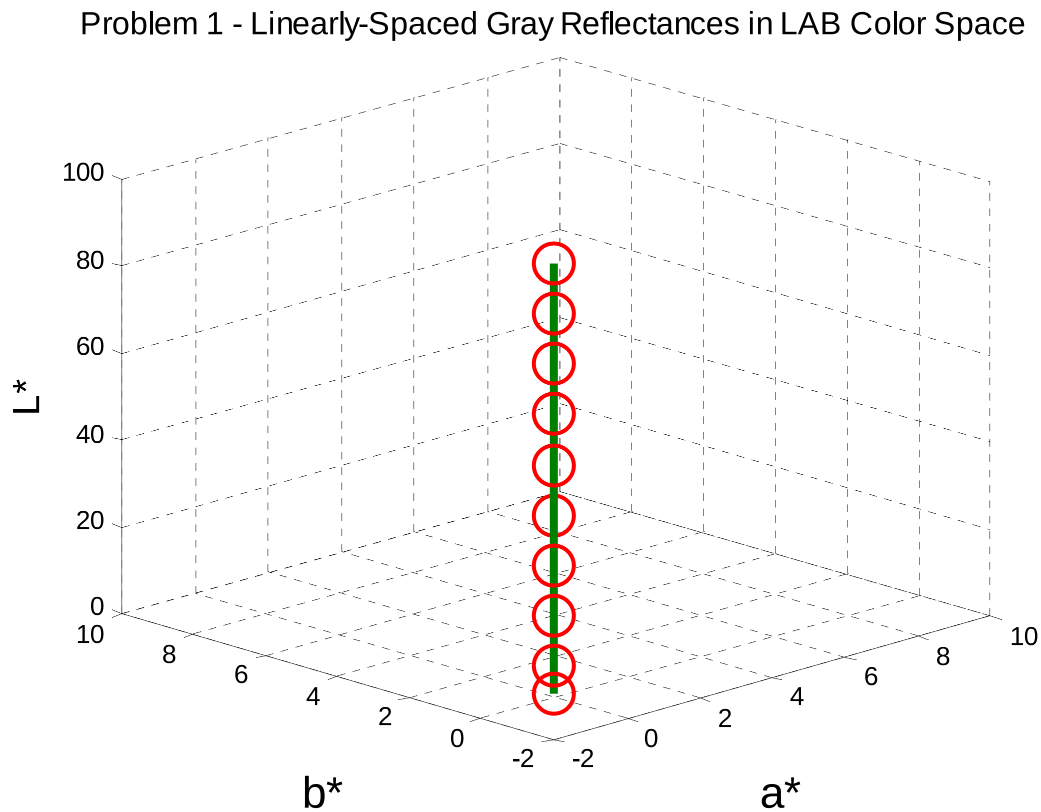
Problem 1 - Linearly-Spaced Gray Reflectances in LAB Color Space
L*
100
80
60
40
20
0
10
8
6
4
2
0
-2
-2
0
2
4
6
8
10
b*
a*

Problem #2 – Visual Sensitivity

$$L^* = 116 \cdot \sqrt[3]{\frac{Y}{Y_W}} - 16$$

$$a^* = 500 \cdot \left(\sqrt[3]{\frac{X}{X_W}} - \sqrt[3]{\frac{Y}{Y_W}} \right)$$

$$b^* = 200 \cdot \left(\sqrt[3]{\frac{Y}{Y_W}} - \sqrt[3]{\frac{Z}{Z_W}} \right)$$

These equations lead to the following variation of L* for linearly spaced reflectances:

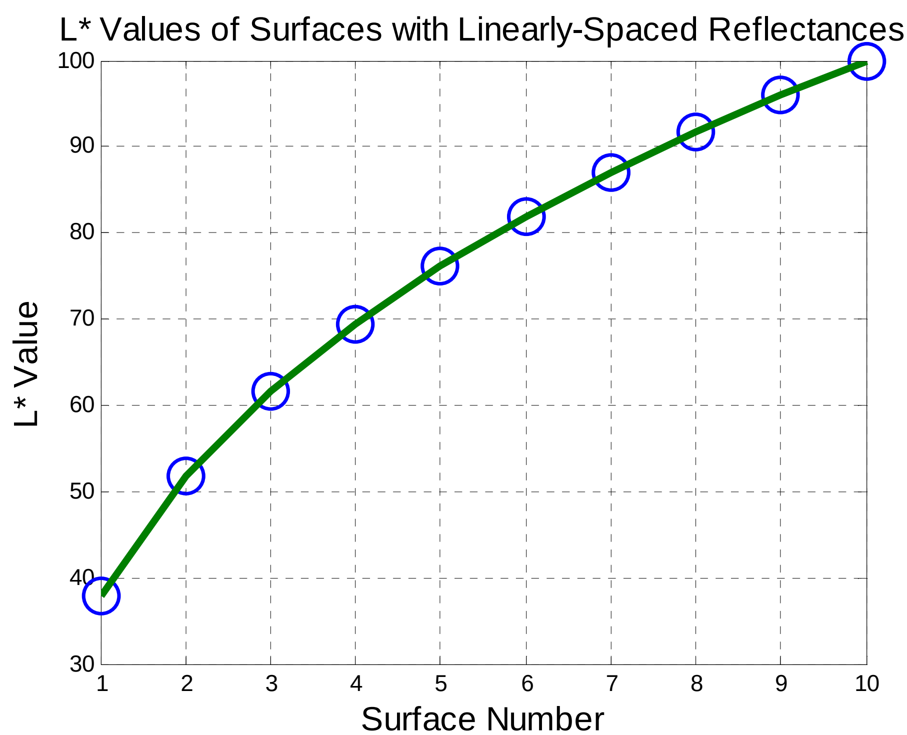

To measure visual sensitivity, we focus on the *distance* between gray level reflectances in L* space; since L* quantifies the discriminability between spatially uniform targets, the eye can better discern reflectances when their L* values are spaced farther apart.

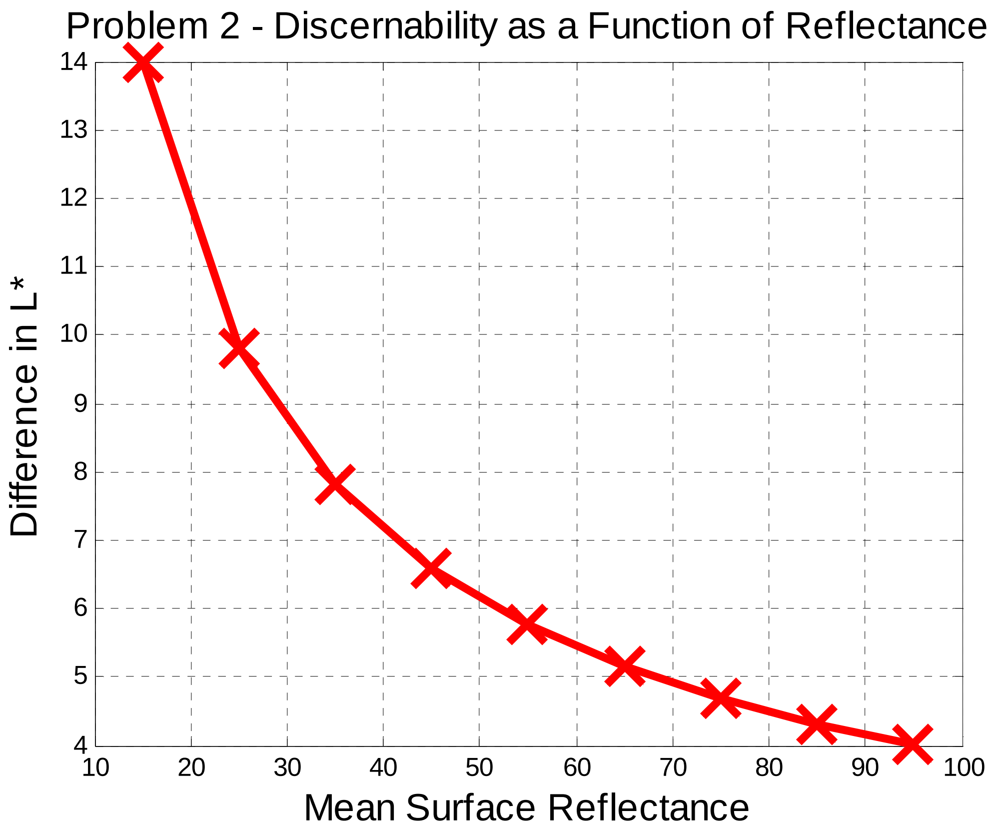

Because

$$L^* = 116 \cdot \sqrt[3]{\frac{Y}{Y_W}} - 16$$

is a concave function of luminance Y, the spacing between L* values of consecutive gray levels decreases at higher surface reflectances. In other words, the human eye more readily distinguishes two targets at low mean surface reflectance than it would at higher reflectances. As a result, our visual sensitivity is greater for distinguishing the mean reflectance of ***dark*** surfaces.

Problem #3 – The CIELAB Metric

Suppose we measure the XYZ values of each of a pair of lights. According to the CIELAB metric, we must also measure the **XYZ values of a reference white surface** before we can predict the discriminability of these lights. Without a reference white to establish a scale against which other reflectances compare, we cannot determine whether the two lights seem far apart or close together in perception. These white values reflect basic intensity sensitivity in the visual system; relative to the other surface reflectances in the system, they inform what normalized linear combinations of red, green, and blue (X, Y, Z) yield pure white (Y = 100) in the gamut of our visual system, and how sensitive our eyes are to contrast about this reference point, answering the question: "How white is white?" In particular, (X_W, Y_W, Z_W) represent the ambient lighting that sets the environment in which our eye views a scene. The instrumentality of these quantities reveals that our visual system is most sensitive to reflectances *relative* to the lighting around them; we perceive a color or intensity only so much as it protrudes from the environment around it. Without juxtaposition, nothing is remarkable. Just as the gray X experiment evinces how important surrounding colors are to our perception of color and pattern, the parameters of white light quantify the importance for our visual system to establish visual differences, discrepancies, or changes in color as a basis for perception. In effect, the background illumination or ambient lighting establishes what is visible and what is not.

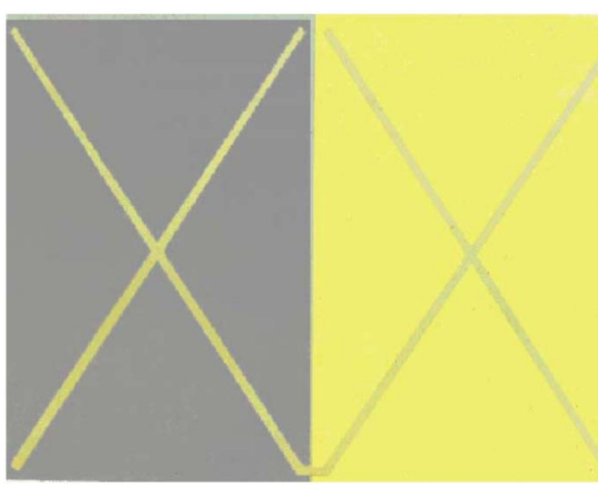

Problem #4 – ΔE Spacing of Five-Unit Perturbations

Reference Point $\mathbf{p} = (\mathbf{X_0}, \mathbf{Y_0}, \mathbf{Z_0})$	Perturbed Point $(\mathbf{X_0} + \Delta\mathbf{X}, \mathbf{Y_0} + \Delta\mathbf{Y}, \mathbf{Z_0} + \Delta\mathbf{Z})$	ΔX	ΔY	ΔZ	ΔE
[50, 100, 100]	[45, 100, 100]	-5	0	0	**13.6955**
[50, 100, 100]	[55, 100, 100]	5	0	0	**12.8104**
[50, 100, 100]	[50, 95, 100]	0	-5	0	**9.3386**
[50, 100, 100]	[50, 105, 100]	0	5	0	**9.0322**
[50, 100, 100]	[50, 100, 95]	0	0	-5	**3.3905**
[50, 100, 100]	[50, 100, 105]	0	0	5	**3.2793**

Apparently, perceived differences to light are most sensitive to X values and least sensitive to Z values for a fixed displacement in XYZ space. This behavior does not surprise us since the $\bar{z}$ tristimulus values are traditionally the highest, making a 5-unit perturbation seem smaller relative to perturbations in X and Y.

Problem #5 – Using CIELAB to Compare Images

Color Coordinate	Yellow	Blue	Green
R	1	0.25	0.625
G	1	0.625	0.8125
B	0	1	0.5
PART (A.) – RGB to XYZ Values			
X	62.0592	52.0468	57.0530
Y	84.6287	58.8485	71.7386
Z	13.7069	127.3905	70.5487
PART (B.) - WHITE POINT [X Y Z] = [95 100 108]			
L*	93.7229	81.2082	87.8426
a*	-39.1022	-9.8722	-25.7505
b*	88.6691	-43.7166	5.5048
PART (B.) - WHITE POINT [X Y Z] = [108 100 95]			
L*	93.7229	81.2082	87.8426
a*	-57.2587	-26.9945	-43.4050
b*	84.2790	-52.9466	-2.0749

The X and Z values of the white point affect only the a* and b* components of the CIELAB values. The value of a* nearly doubles for yellow and green and nearly triples for blue. Meanwhile, b* slightly decreases for yellow and decrement by nearly 10 units for blue and green. On the other hand, the L* component remains unchanged since the luminance Y is still 100. All in all, these alterations reflect the strong solitary dependence of L* on luminance Y, and the looser correspondences between a*~X and b*~Z, although both a* and b* depend on Y as well.

PART (C.) - ΔE Color Differences

$$\Delta E_{YB} = 136.1505$$

$$\Delta E_{YG} = 84.4343$$

$$\Delta E_{BG} = 52.1429$$

According to the ΔE metric, yellow and blue are the most dissimilar. Yellow and green are also dissimilar, while blue and green are the closest together of the three in terms of uniform color perception.

PART (D.) - ΔE Color Differences between Images

$$\Delta E_{YB1-G} = 68.2886$$

$$\Delta E_{YB2-G} = 68.2886$$

Thus, according to the ΔE metric, both the high-frequency striped image imYB1 and the thickly striped image imYB2 are equally similar to the uniformly green image imG. In other words, the human visual system should perceive roughly the same distance (or dissimilarity) between the green image and each of the two striped images.

PART (E.) – Improving the CIELAB Color Metric

However, despite the equality in ΔE, the more frequently striped image looks closer to the uniform green image in my eyes than the thickly striped image, most likely because my eyes essentially lowpass filter the image, blurring together the thick stripes in an effort to digest the whole of the picture.

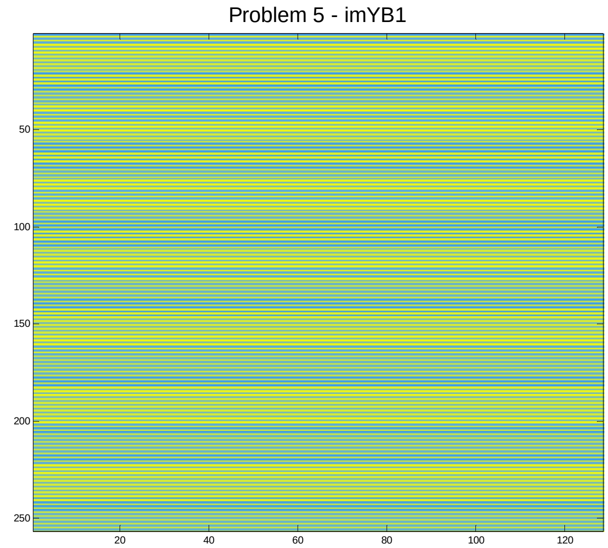

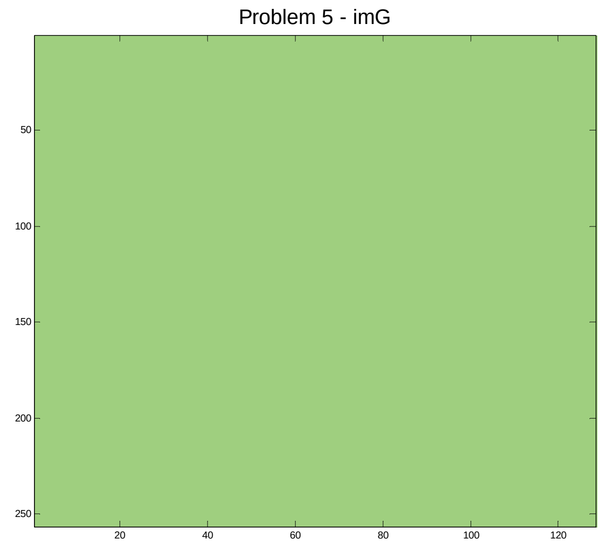

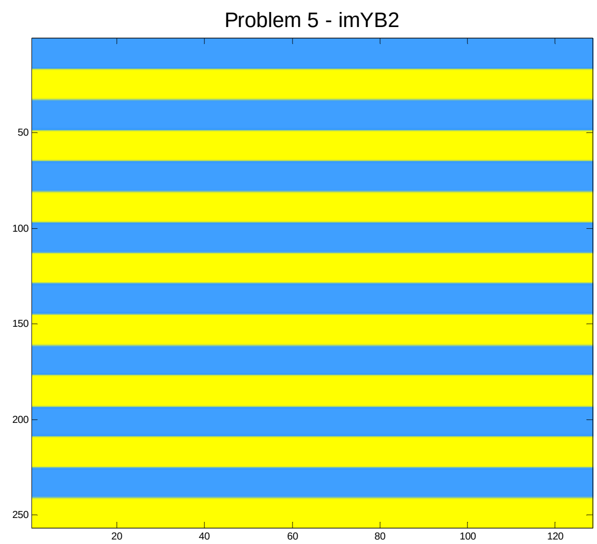

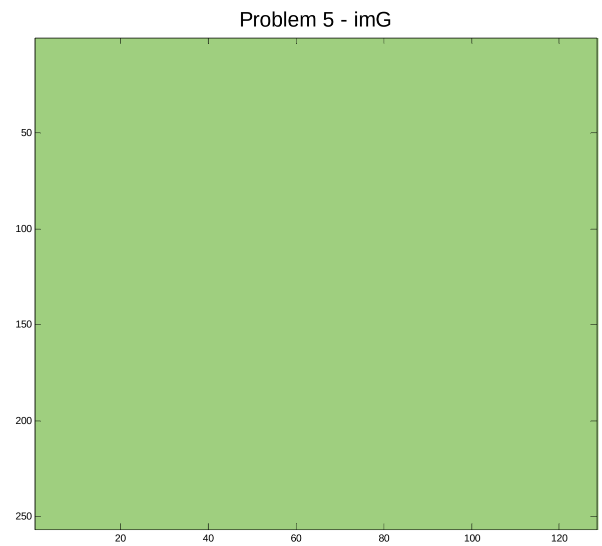

The stripes in imYB1 are much finer and alternated much more frequently than the stripes in imYB2, so, when both images are juxtaposed beside the constant green field of imG, imYB1 actually resembles the uniform green field much more closely than imYB2 from the human eye's perspective. Our visual sensitivity to low spatial frequencies causes us to perceive the thickly striped image imYB2 extremely differently from the uniform green field of imG, whereas, at the high spatial frequency of imYB1, our eye tends to digest the big picture, blurring together the thin stripes even though we see them; despite our knowledge of the high frequency, our eyes betray us as they lowpass filter the fine detail, allowing the general uniformity of green in the union of all the yellow and blue stripes to protrude forth most prominently. All in all, our eye seems to join the stripes, as if blurring them together, when judging the overall image; even armed with the knowledge that imYB1 has frequently varying stripes, our eye relentlessly merges information from the rapidly alternating blue and yellow to form a general green background that appears much closer to the green field than thick, slowly varying stripes.

However, the CIELAB metric, designed to compare visual perception of uniform fields, treats the entirety of imYB1 and imYB2 identically. Averaging the ΔE values across the entire image, the CIELAB metric extracts the same information content from imYB1 and imYB2, which contain the same number of yellow pixels and blue pixels. Thus, from a frequency-blind uniform-field perspective, imYB1 and imYB2 are identical, so their CIELAB values and hence their distance from the green field match perfectly. In other words, CIELAB fails to incorporate the frequency of the pattern and the local distribution of those color pixels throughout the image into the metric, so the distance ΔE that we computed actually neglects the frequency-sensitivity of our eye as we differentiate stripes or variations of different frequency. Because our eye is quite sensitive to local variations and low-frequency patterns, this shortcoming results in discrepancies between visual discriminability and the calculated ΔE values, as explained in the work and writing of Poirson & Wandell (1993) and Kelly (1996).

To improve the CIELAB metric to consider the eye's frequency sensitivity, we might separate the original color image into opponent colors (such as yellow and blue here, or red and green elsewhere) to account for the blending that our eyes perform. Most importantly, however, we can subsequently apply a spatial lowpass filter to the opponent colors images, paralleling the function of our eye as we attenuate higher frequencies prior to computing the CIELAB values. Much as our eye absorbs the larger picture of a high-frequency pattern, our spatial filter would extract and accentuate the low frequencies clearly visible and readily distinguishable while downplaying the effect of higher frequencies such as the stripes seen in imYB1. As a result, the subsequent conversion of the filtered image into CIELAB space would give greatest weight to the low frequencies (thick stripes), mirroring the sensitivity of our eye to lower-frequency patterns.

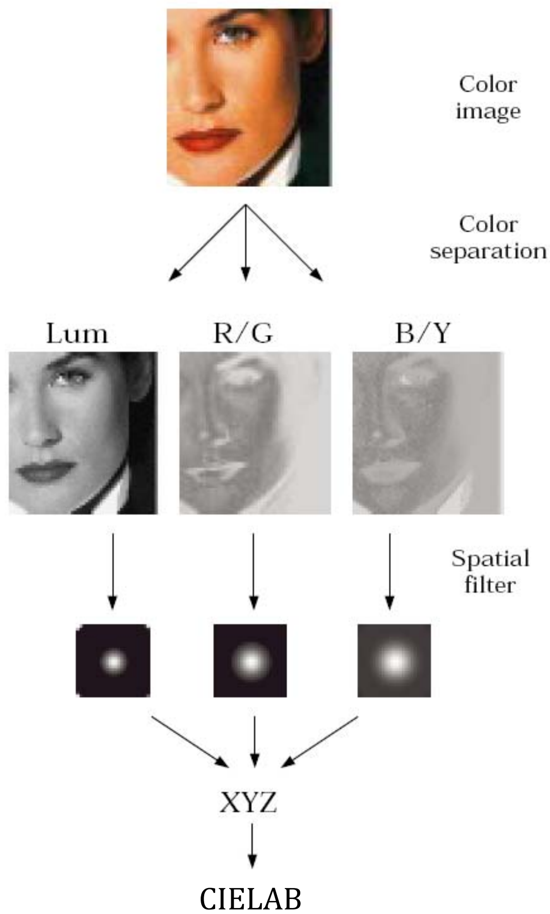

For example, if we applied this method to our current problem, then the lowpass filter for the yellow-blue opponent colors image would smoothen the high-frequency variation in imYB1 and blend together the yellow and blue stripes to form green even before CIELAB values are computed, mimicking our eyes' tendencies to blur and extract slowly-varying, overarching patterns. Essentially, just as we interpret imYB1, the lowpass-filtered version would look much like a uniform green field, allowing its CIELAB conversion to be close to the CIELAB conversion of the green field, thereby decreasing the ΔE metric to a small quantity true to the similarity our eyes perceive. Meanwhile, the thickly-striped low-frequency imYB2 image would undergo little change before CIELAB conversion, since it contains no high spatial frequencies beyond the cutoff of our spatial filter, so its CIELAB and ΔE values would remain the same. However, by removing the high frequencies of the first image imYB1, we have effectively brought imYB1 closer to imG in CIELAB space, accomplishing our goal of reflecting visual similarity in our ΔE metric.

Xuemei Zhang and Brian Wandell devised the S-CIELAB model to extend the similarity metric into the frequency domain, and many of these suggestions (as well as the block diagram) derive from their filter bank approach, as detailed in "A Spatial Extension of CIELAB for Digital Color Image Reproduction" from *Society for Information Display* (1996):

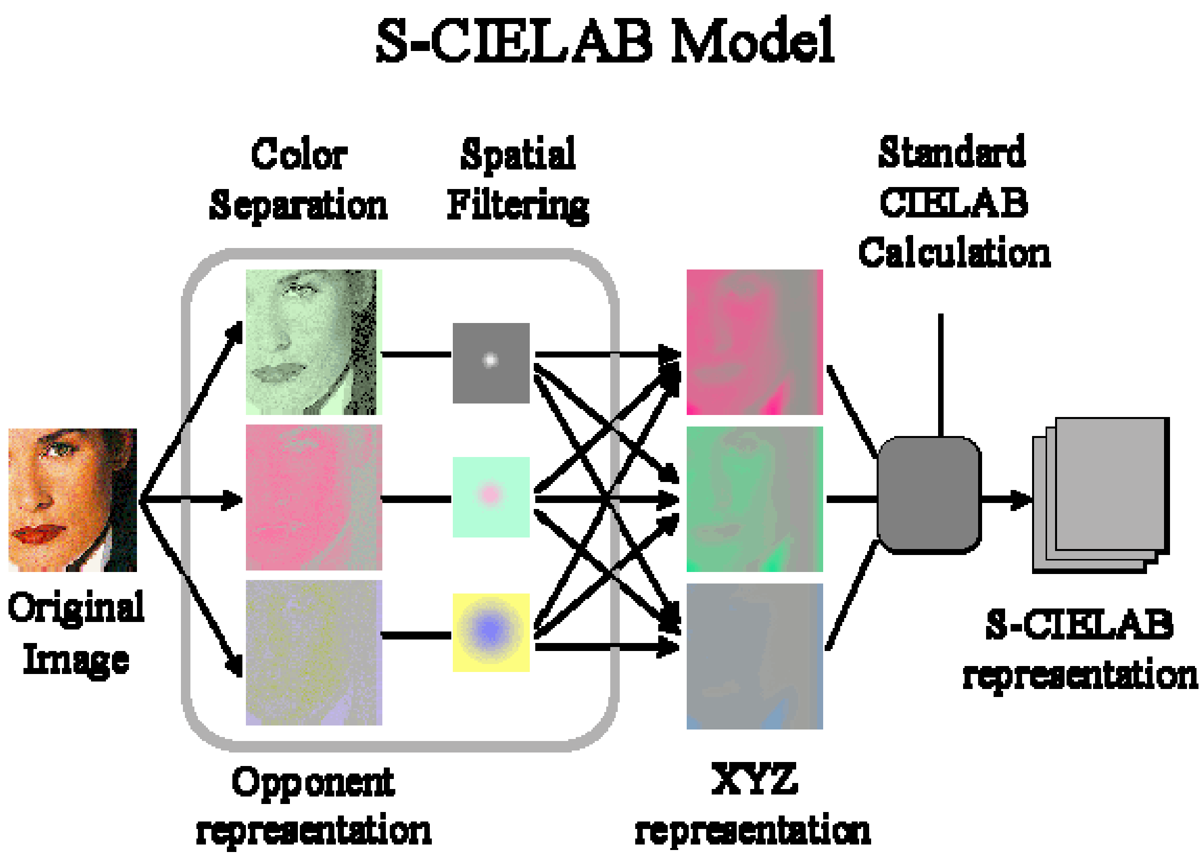

http://white.stanford.edu/~brian/scielab/introduction.html

http://white.stanford.edu/~brian/papers/ise/scielab/scielab.pdf