A Study of Scene Structure in the Saturation Component of Color Images

Bruce A. Thomas, Robin N. Strickland
Department of Electrical and Computer Engineering
University of Arizona, Tucson, Arizona 85721

Abstract

A tenet of a new class of color image enhancement algorithms involves the observation that the saturation component of color images often contains what appears to be valid image structure depicting the underlying scene. In this work we present the findings of a study of the structural correspondence between the saturation and luminance components of a large database of color irnages. Various statistical relationships are identified. The correspondence of edges at different scales in the sense of Marr's theory of vision is also observed. Several new color image enhancement algorithms which exploit these unique characteristics are described.

Keywords: Color image processing, HSL color space, image enhancement, human color perception

1. Introduction

With the advent of low-cost digital color input devices in recent years, the realm of digital color imaging is rapidly taking hold. This is particularly evident at the consumer product level where digital color cameras and scanners are entering the market in increasing numbers.[1,10] Unfortunately, the field of digital color image processing remains fairly nascent, with the bulk of current algorithms being rooted in component-separable and luminance-based processing strategies. The unique nature of color image information offers many opportunities for new color image enhancement and restoration techniques.

Upon examining the hue, saturation, and luminance (HSL) components of the HSL description of a color image, it is common to find that the luminance and saturation components bear a remarkable similarity to one another. Several examples of this are shown in Fig. 1. The complementary nature of the luminance and saturation information invites new strategies for performing image enhancement on color images, ones which exploit the relevant structural information contained in the saturation component. Several investigators have undertaken work based on this strategy.[12,13]

1.1 Background

The fundamental basis for the modern description of color information lies with the trichromatic color theory first postulated by George Palmer in 1777.[8] Palmer proposed, and it has since been verified, that there are 3 key receptors for color information in the human visual system. As such, color information can be uniquely represented by a set of 3 tristimulus values, the most familiar being the RGB description used in color television. A variety of alternative 3-dimensional descriptions have also been derived. Some of these form color spaces geared for color difference measurement (e.g. CIE $L^*u^*v^*$) while others bear more perceptually relevant coordinates, e.g. HSL. Indeed, the HSL (and its counterpart the HSV) color space[4,11] has components that relate directly to hue (base color),

Figure 1. Examples of the luminance (top row) and saturation components (bottom row) of color images (the saturation components have been negatively scaled using 1 - 5(x, y)).

saturation (color purity), and lightness/value (brightness). This close correspondence to the psychophysical aspects of human color perception have made the HSL color space, and variants thereof, a desirable choice for color image segmentation and enhancement schemes.[2,3,9,12]

From a mathematical viewpoint, the HSL color space is essentially a cylindrical representation of the rectangular RGB color space (see Fig. 2). The main diagonal of the RGB color cube forms the (unbiased) luminance axis having zero luminance value at the origin ($R = G = B = 0$) and maximum luminance value at its apex ($R = G = B = 1.0$). Planes of constant luminance lie perpendicular to this axis. Hue and saturation are determined within these planes as polar coordinates extending from an origin defined by the point of intersection between the luminance axis and each plane of constant luminance. As shown in Fig. 2, hue is defined as the angle subtended in the plane with respect to a common reference. Saturation is defined as the radial length in the plane relative to the longest radial dimension allowed for a given polar angle. These form the coordinates of the "unbiased HSL triangle" model as reported by Smith[11] in 1978. The generalized version of this model developed by Levkowitz and Herman[4] was used for this report (see reference [4] for complete details).

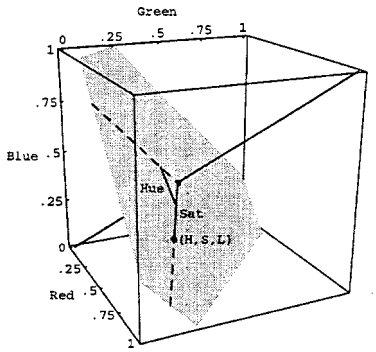

Figure 2. Graphical depiction of the HSL color space.

During the mid-1980s, use of the HSL color space and early observations of the similarity between luminance and saturation components prompted the development of a new paradigm for color image enhancement.[12] In this approach, high-frequency spatial information from the saturation component is fed back into the luminance component as a means of enhancing color image contrast. This technique can serve to bring out details of an image that have low luminance contrast. More recently, others have seized upon this paradigm. Toet[13] has performed a pilot study of a multiscale approach wherein information from the saturation component is fed back into the luminance component using a nonlinear pyramid recombination. Similarly, Liu and Yan[5] have proposed a related scheme that employs a color difference measure rather than information from the saturation component directly. Thus a growing number of techniques are being molded after the saturation-feedback paradigm, each a unique species of this new class of color image enhancement algorithms.

1.2 Goals

The aim of the present work is to investigate the nature and extent of the unique correspondences that tend to exist between the luminance and saturation components of color images. This is carried out by means of an empirical study of the relationships that exist between the saturation and luminance components of images from a large database of color images. The results are intended to fuel interest in the saturation-feedback paradigm for color image enhancement, as well as to incite new thinking towards related color image enhancement schemes.

2. Experimental Observations

A database of 35 digital color images was established for the purposes of this report. All but three of the images came from Kodak PhotoCDs: 17 from a Kodak PhotoCD sampler (part no. 151010-01) and 15 from a locally created PhotoCD (electronic transfers of various Ektachrome and Kodachrome slides). The remaining images are standard test color images available in the public domain. The color images were chosen to span a variety of scenes including those of people (individuals and crowds), outdoor settings (ones containing man-made structures), indoor settings, landscape (natural settings), and objects (e.g. boats, statues etc.). These images were then subjected to a variety of experiments in an effort to answer some basic questions about the structural correspondence between the saturation and luminance components of color images. In each case, the saturation and luminance components were extracted, normalized to an identical range of values, then treated as independent grayscale images. In the data summaries to follow, the data are grouped according to the following image categories:

Categorization Label	Description
P	People
S	(outdoor) Scene
I	indoor (Scene)
N	Natural (landscape)
O	Objects

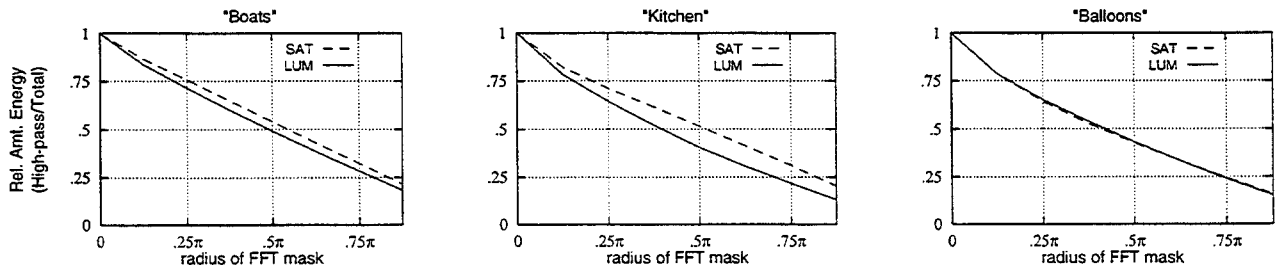

Figure 3. Plots of the cumulative amounts of high-pass spectral energy for the saturation and luminance components of the images in Fig. 1.

2.1 High Frequency Content

Which component typically contains more high-frequency spatial structure? In 1987 Strickland[12] claimed that the saturation component of color images typically contains higher spatial frequency content than its luminance counterpart. In order to further substantiate this, we have examined the relative amounts of high-frequency spectral energy in the FFTs of the luminance and saturation components of the images in our database. A circular FFT mask, centered about DC, was used to isolate the high-frequency portions of the respective FFTs. The sum of the absolute energies of the isolated high-frequencies was then measured relative to the total energy (sans DC) of the FFT. By doing this with masks of varying radius, it is possible to examine the cumulative distribution of the relative amounts of high-frequency spectral energies. This was done for both the luminance and saturation components of the images in our database. Fig. 3 shows the distributions resulting for the images of Fig. 1. It should be noted that the saturation components generally seem to exhibit more high-frequency energy. This fact can be borne out by examining the ratio of saturation-to-luminance high-pass energies. Ratios of this sort were taken for each of the images in our database. All but 2 of the images showed ratios exceeding 1 which indicates that the saturation component contains more high-pass spectral energy than its luminance counterpart. A summary plot of the computed ratios is given in Fig. 9a.

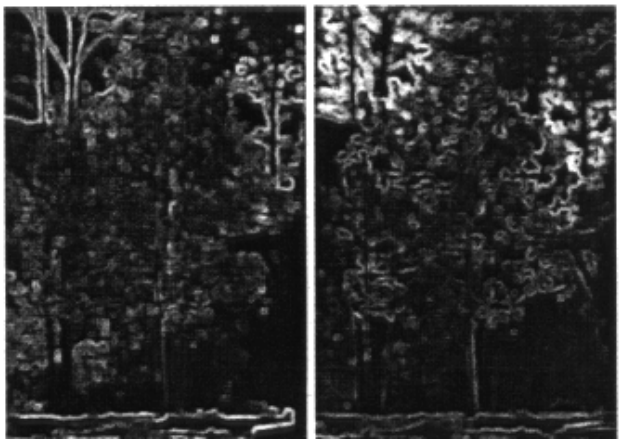

Figure 4. Local variances maps for image of a natural scene (a) luminance (b) saturation.

2.2 Regions of Activity

Do the luminance and saturation components tend to exhibit "activity" in the same local regions? To answer this question local variance was computed over 5×5 pixel neighborhoods as a measure of "local activity." Fig. 4 shows a typical example of the *local variance maps* so produced (black represents zero variance, white represents maximum variance). It is interesting to observe that the saturation component exhibits more activity along certain edges than does the luminance component. The degree of concurrence of "local activity between the saturation and luminance components was measured by first thresholding the individual variance maps at their 40th percentile levels. A logical AND operation was then used to assess the degree of overlap between the resulting "most active" regions. In order to avoid false measurements due to numerical instabilities, pixels lying in the 5×5 neighborhoods of extreme luminance values were ignored (the computation of saturation becomes particularly unstable for small and large values of luminance). Using these methods, we measured the percentage overlap in the (40%) most active pixel locations of the saturation and luminance components of the images in our database. A scatter plot summary of the resulting data is given in Fig. 9b.

2.3 Local Correlatedness

How well do the luminance and saturation components of a color image correlate? For the answer to this we used the centered value of a local correlation coefficient product to assess the degree of local correlation. The product was computed over a 5×5 pixel window that was registered over identical locations in the saturation and luminance component images. Measurement of the local correlatedness thus involved computing the following product at each pixel location

$$LCC(x_o, y_o) = \frac{\sum_{x,y \in W}\left[L(x,y) - \bar{L}_W\right]\left[S(x,y) - \bar{S}_W\right]}{\left\{\sum_{x,y \in W}\left[L(x,y) - \bar{L}_W\right]^2\right\}^{\frac{1}{2}}\left\{\sum_{x,y \in W}\left[S(x,y) - \bar{S}_W\right]^2\right\}^{\frac{1}{2}}}. \quad (1)$$

This product delivers values lying in the interval [–1,1]. The set W identifies the region of support for a 5×5 window centered at pixel location (x_o,y_o) and L_W and S_W represent the mean value of the luminance and saturation components, respectively, as computed over the specified region of support. The resulting *local correlation coefficient map* was then thresholded to reveal those pixel locations for which the degree of correlation was very strong, i.e.. greater than 80% positive (shown in white) or 80% negative (shown in gray). A sample output from this thresholding scheme is given in Fig. 5. This image, which is from a standard Kodak PhotoCD sampler, displays a significant degree of local correlatedness between its saturation and luminance components.

The lighthouse image used in Fig. 5 was not the only color image to exhibit strong local correlatedness between its luminance and saturation components. Several other images exhibited even greater amounts. In an effort to quantify this for each of the images in our database, we thresholded each image's *local correlation coefficicnt map* using the ±80 scheme described above and counted the percentage number of thresholded pixels. This value was taken to be the correlatedness for that image and describes the extent to which the image experiences strong correlations between its luminance and saturation components. Note that pixel locations bearing zero variance (in either the luminance or saturation component) were excluded from the local correlatedness computations as the correlation coefficient becomes unbounded at these locations. We used the above process to measure the local correlatedness of each of the images in our database. These data are summarized in the scatter plots of Fig. 10.

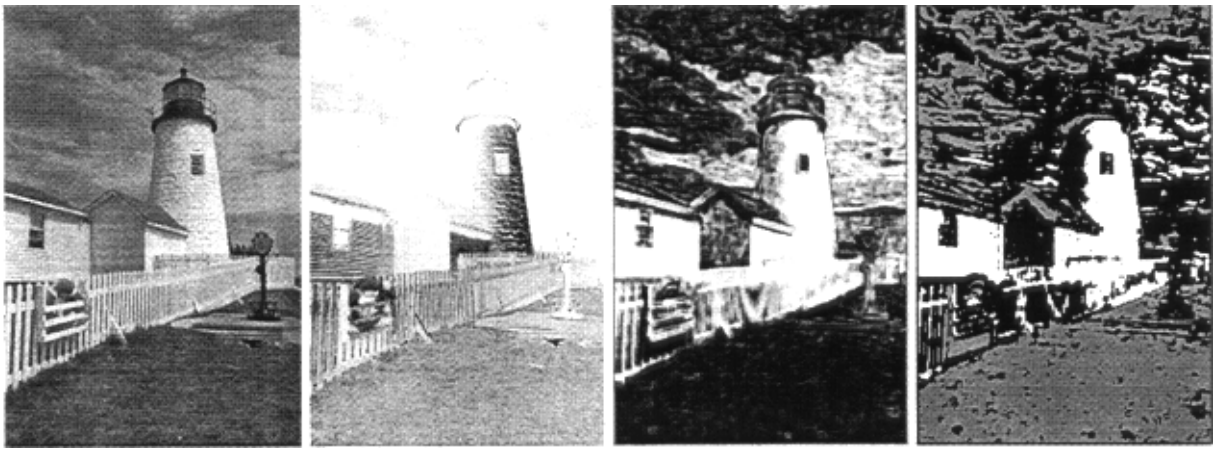

Figure 5. Demonstration of a local relatedness analysis (a) LUM (b) 1-SAT (c) Local correlation coefficient map (d) Map thresholded at ±80 extrema.

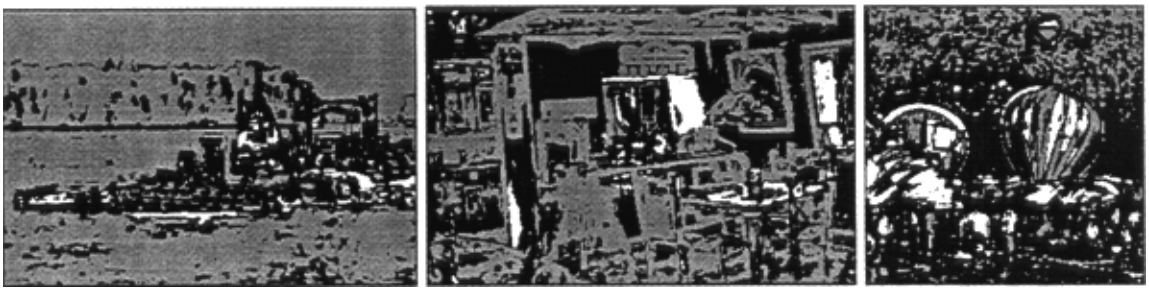

Figure 6. Examples of LUM-SAT correlatedness (a) strong [70%] (b) medium [47%] (c) weak [26%].

Figure 7. Examples of the correspondence between SAT (black) and LUM (gray) edge maps derived from two of the images of Fig. 1 using $\sigma = 2.934$.

The degree of local correlatedness experienced by the images in our database spanned a broad range of values. Some examples of the different degrees to which the luminance and saturation components exhibit strong correlations are given in Fig. 6. Here we show the positively and negatively thresholded versions of the *local correlation coefficient maps* for color images exhibiting strong, medium, and weak amounts of correlatedness.

2.4 Edge Map Relationships

What is the structural correspondence, at varying scales, between the luminance and saturation components of color images? For this study, a Marr-Hildreth[6] edge detection scheme was used to derive edge maps from the luminance and saturation components using Laplacian of Gaussian (LOG) filters of varying size. These form the basis for what Marr[7] called the *primal sketch* which, according to his computational theory of vision, can be used to generate a description of the underlying scene. For our purposes, LOG filters of size σ = 0.978, 1.956, and 2.934 were used to generate edge maps at three roughly octave-separated scales. These filter dimensions were chosen to coincide with those obtained by expanding the first three levels of a Laplacian pyramid to full (original) image size.

The edge maps derived from the saturation and luminance components of the images in our database tended to display unique and curiously overlapping representations of the underlying scene. Two examples of this are shown in Fig. 7 where we have superimposed the saturation and luminance component edge maps from a common scale of

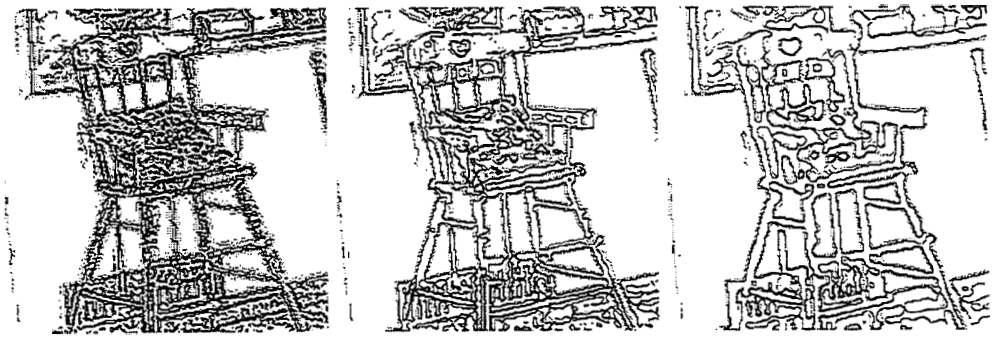

Figure 8. Close-up view of the correspondence of SAT and LUM edge maps at different scales (a) $\sigma = 0.978$ (b) $\sigma = 1.956$ (c) $\sigma = 2.934$.

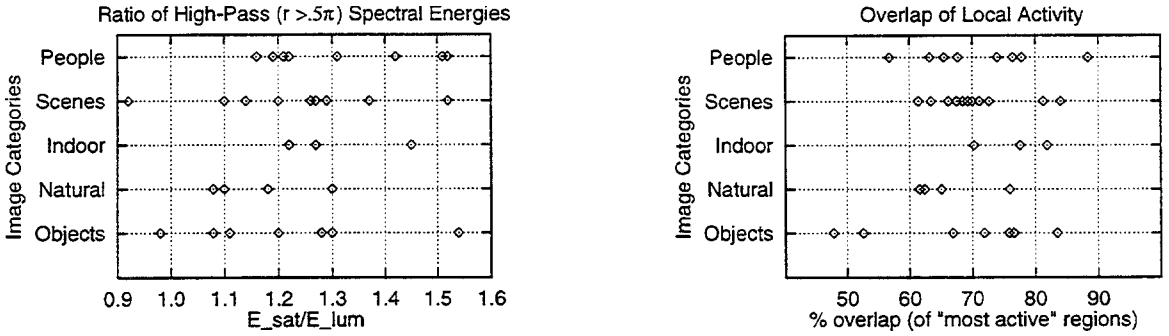

Figure 9. Summaries of saturation vs. luminance structural comparison data (a) high-frequency content (b) regions of local activity.

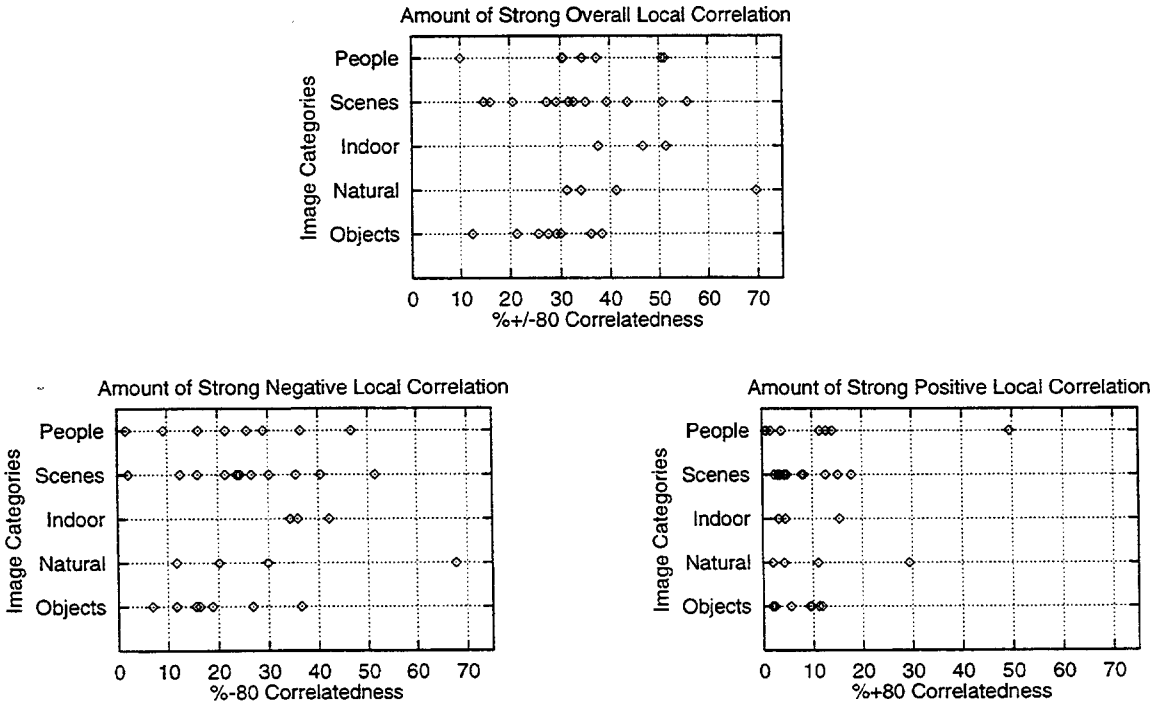

Figure 10. Summaries of saturation vs. luminance local correlatedness data (a) net correlatedness (b) negative correlatedness (c) positive correlatedness.

analysis. The edges derived from the luminance component are shown in gray and the edges from the saturation component are shown in black. It is interesting to observe the complementary fashion with which the two sets of edges describe the scene. Fig. 8 gives an example of the different scales of edge information obtained from this analysis.

The amount of edge information at the various scales of the saturation and luminance components varied from image to image. In an effort to quantify this, the ratio of the number of luminance edge pixels to the number of saturation edge pixels was tabulated at each scale. Any pixel lying in the 3×3 vicinity of a significant sign change in the LOG

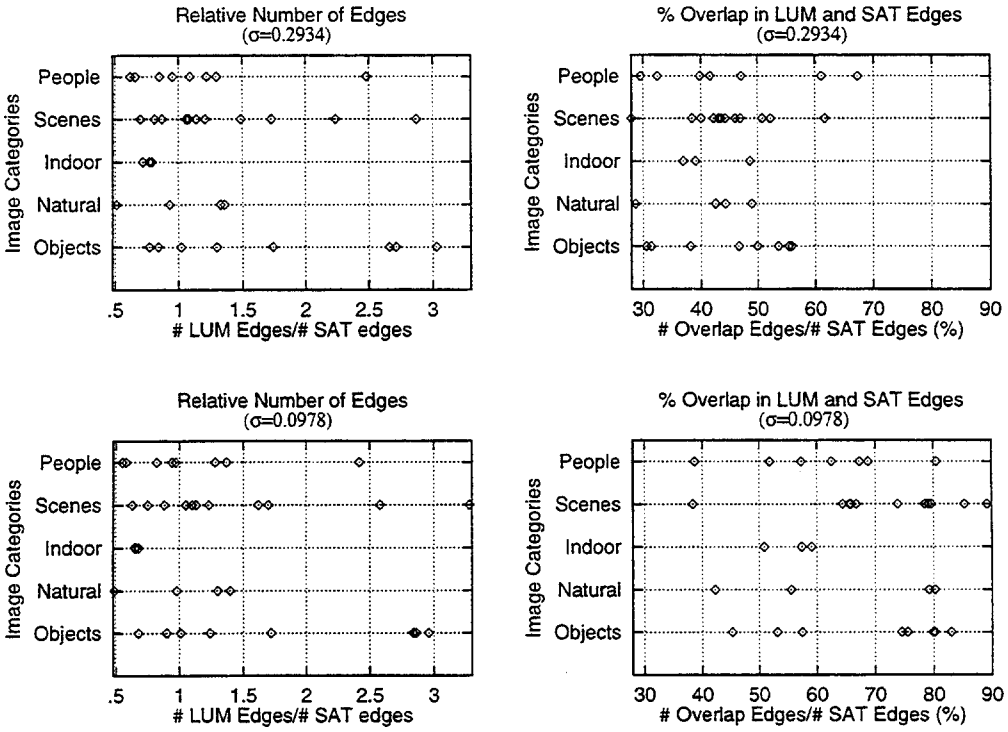

Figure 11. Sundries of saturation vs. lununance edge structure comparisons.

filter output was deemed to be an edge pixel. A mask corresponding to areas of zero or near-zero variance was also used to screen out false edges associated with areas of low variance. In addition, the degree of overlap between the luminance and saturation edge maps was assessed by computing the percentage overlap experienced by each of the individual edge maps. Graphical summaries of these data are presented in Fig. 11.

3. Analysis of Results

The data obtained from the experiments of section 2 have been summarized in the form of scatter plots. These are shown in Figs. 9, 10 and 11. In each, the data has been separated according to the image categories described at the beginning of section 2. This was done to reveal any trends that may be related to specific categories of images.

3.1 High Frequency Content

By and large, the saturation component of color images seems to contain more high-frequency spatial content than its luminance counterpart. This can be seen in Fig. 9a where the ratios of saturation-to-luminance high-pass energies are predominantly greater than 1. Only two of the images in our database countered this trend, and even then only slightly. Of note are the cases where the saturation component houses 40-50% more high-frequency spatial structure than its luminance counterpart. These are images that are probably well-suited for enhancement by a saturation-feedback scheme.

3.2 Regions of Activity

In general, there seems to be remarkably good but not complete agreement between the locally "most active" regions of the luminance and saturation components of color images. As can be seen in Fig. 9b, the bulk of the images in our database show 60-80% overlap. This hints to the existence of a strong amount of structural coherence between the luminance and saturation components.

3.3 Local Correlatedness

All of the images in our database showed areas where the degree of correlation between the luminance and saturation components was very strong (> +80% or < -80%). These are regions where the scene structure conveyed by the luminance and saturation components tend to corroborate. As can be seen in the top portion of Fig. 10, a majority of the images exhibited a significant (>30%) amount of strongly correlated pixel locations. Thus, in over 30% of the pixels in these images it may be possible to use information from the saturation component to reinforce the luminance component signal. This can be particularly advantageous in areas of weak luminance but strong saturation variation (note that our measure of correlatedness was based on a correlation coefficient product which is invariant to signal scaling). Areas lying outside the regions of strong correlation may also benefit from such a scheme, but the correspondence of base image structure is not as powerfully guaranteed (which may have its own benefits).

The regions of strong correlatedness can also be broken down into their specific polarities as shown in the bottom portion of Fig. 10. The most noticeable characteristic here is that negative correlations make up the dominant portion of the overall correlatedness. Indeed, it is more common to find strong negative correlations between saturation and luminance than it is to find strong positive correlations. Positive correlations between the saturation and luminance

components seem to occur mostly in areas of brighter luminance, which may serve as a useful gating function for determination of saturation-feedback polarity.

3.4 Edge Map Relationships

The edge maps derived from the luminance and saturation components of the color images in our database provided distinct portrayals of the underlying scene. Among the images tested, there was a noticeable degree of overlap between luminance and saturation component edge maps, but in each case the saturation component maintained independent and structurally viable edge information. The existence of this additional edge information is certainly of use to the enhancement of color images. Some cautions are in order though as sometimes unique but potentially undesireable saturation edges appear along the boundaries of highlights (see the "balloons" images of Figs. 1 and 7).

As can be seen in Fig. 11, there seems to be a tendency for the luminance component to exhibit a larger number of edges than its saturation counterpart. There is also some evidence that for indoor scenes, the saturation component consistently holds a larger number of edges. This could be due to the more diffuse illumination typically found in indoor environments.

An examination of the edge structure at different scales reveals little change in the relative number of edges found. Indeed, the data are nearly identically distributed across all scales investigated (see left-hand plots of Fig. 11). The percentage of unique saturation edges however tended to increase with scale. This is evidenced in the right-hand plots of Fig. 11 where the percentage of overlapping edges diminishes at increased scales. This trend was experienced across all three scales investigated and could be useful to multiscale saturation-feedback schemes for color image enhancement. The evolution of unique saturation information as scale is increased may be linked to the curious tendency for edges in the saturation component to lie adjacent but not necessarily in direct register with those in the luminance component. The causes of this phenomemon are still under investigation.

4. Applications

The saturation component of color images typically contains more high-frequency spatial content than its luminance counterpart. It also bears image structure that concurs to a large extent with that which is contained in the luminance component. These characteristics may be exploited to produce new methods of color image enhancement.

Adaptive enhancement methodologies lie at the heart of one set of possibilities. A simple approach is to use local variance information from the saturation component to identify areas where adaptive contrast emphasis should be applied in the luminance component. This may serve to bring out details that are not immediately evident in the luminance component alone. An alternate approach uses local variance and correlation coefficient information to localize the application of pre-existing saturation-feedback enhancement schemes such as the one proposed by Strickland et. al. in 1987.[12] This would serve to contain the enhancement effects to strictly those regions bearing a certain degree of structural coherence between their saturation and luminance components. Indeed, the strength of the luminance-to-saturation correlation coefficient could also serve as an adaptive weighting function at each pixel location.

Data fusion techniques may lead to another set of possibilities for improved color image enhancement. For images bearing significant degrees of saturation-to-luminance correlatedness, it is appropriate to fuse some of the high-frequency spatial information from the saturation component directly into the luminance component. This could be accomplished with any one of the growing number of data fusion techniques appearing in the literature. Again, as before, this could be done globally or under adaptive control as specified by the local variance or correlation coefficient of the saturation (and luminance) components. These are just a few of the novel ways in which the additional information contained in the saturation component may be exploited to achieve improved enhancement of color images.

5. References

1. K. Balleisen and K. Ryer, "Apple's imaging future,", *MacWEEK*, Vol. **9**, No. 46, pp. 1, 106 (1995).
2. W. J. Carper, T. M. Lillesand, and R. W. Kiefer, "The use of intensity-hue-saturation transformations for merging SPOT panchromatic and multispectral image data," *Photogram. Eng. Remote Sensing*, Vol. **56**, pp. 459-467 (1990).
3. J. Y. Kim and Y. H. Ha, "Color image enhancement using color constancy based on modified IHS coordinate system," in *Intelligent Robots and Computer Vision XII*, David P. Casasent, Editor, *Proc. SPIE* **2055**, pp. 359-364 (1993).
4. H. Levkowitz and G. Herman, "GLHS: A generalized lightness, hue, and saturation color model," *CVGIP: Graphical Models, and Image Processing*, Vol. **55**, No. 4, pp. 271-285 (1993).
5. N. Liu and H. Yan, "Improved method for color image enhancement base on luminance and color contrast", *Journal of Electronic Imaging* Vol. **3**, No. 2, pp. 190-197 (1994).
6. D. Marr and E. Hildreth, "Theory of edge detection,", *Proc. R. Soc. London*, Vol. **207**, pp. 187-217 (1980).
7. D. Marr, *Vision*, W. H. Freeman and Company, San Francisco, 1982.
8. G. Palmer, *Theory of Colors and Vision*, Leacroft, London, 1777.
9. F. Perez and C. Koch, "Toward color image segmentation in analog VLSI: algorithm and hardware," *International Journal of Computer Vision*, Vol. **12**, No. 1, pp. 17-42 (1994).
10. K. Ryer, "Kodak and Fuji focus efforts on digital cameras *MacWEEK*, Vol. **10**, No. 1, pp. 1, 142 (1996).
11. A. R. Smith, "Color gamut transform pairs," *Computer Graphics*, Vol. **12**, No. 3, pp. 305-325, 1978.
12. R. N. Strickland, C. S. Kim, and W. F. McDonnell, "Digital color image enhancement based on the saturation component," *Optical Engineering*, Vol. **26**, No. 7, pp. 609-616 (1987).
13. A. Toet, "Multiscale color image enhancement," *Pattern Recognition Letters*, Vol. **13**, pp. 167-174 (1992).

published previously in SPIE, Vol. 2657, page 32