Isoluminance: a Color Technique for Visualizing Multivariable Medical Image Data

BM Hemminger

Department of Radiology
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-7510

ABSTRACT

We have developed a technique for visualizing medical image data sets that have multiple values at each physical location. These data sets are increasingly common as physicians attempt to correlate between modalities (for instance, CT and nuclear medicine, MR and PET, CT and MR) as well as within modalities (for instance, MR metabolic and anatomical scans). Our technique, termed "isoluminance", is designed primarily for displaying data sets that have two scalar values associated with each physical (x,y) location on a two dimensional scan. A perceptually uniform luminance scale is used to encode one dimension. At each step of the luminance scale a set of isoluminant hue steps are used to encode the other dimension. The hue scale is chosen to be perceptually uniform and as "natural" as possible. The resulting data set can then be displayed as a single image on a color display. We have found observers using our technique are able to comprehend both of the data sets, to understand relationships between the data sets, and, when using interactive manipulation techniques, are able to select or label specific features of the data set.

We will report on the development of the isoluminance model from visual and perceptual tenets. We will discuss the choice of mapping the data to luminance and hue attributes, the selection of hue and luminance scales, the advantages and disadvantages of this technique applied to common multivariable visualization problems, and the implementation of this technique on a standard workstation. Finally, we will discuss our experience in applying the technique to several clinical problems.

1. INTRODUCTION

An important problem in medical image display is how to correlate information in images from multiple acquisitions that are spatially registered. In radiology it is becoming increasing common to compare images from multiple modalities, e.g. CT, MRI, PET and nuclear medicine, as well as intra-modality, such as different MRI acquisitions (i.e. T1, T2, different spin echoes, MR Spectroscopy (MRS), etc). These problems are representative of the general problem of understanding large multivariable data sets, such as encountered in geosciences, cartography, computational fluid dynamics, finite element analysis, statistics, weather modeling, computer simulations, etc.[1] This area is called *scientific visualization* in the computer graphics community.[2]

Scientific visualization involves selecting information from data sets and conveying it to the user through a sensory channel (usually visual). I will use the common term, *data visualization*, to refer to the process of the human observer visually perceiving information from a computer display screen. I will use the term, *data reduction*, to refer to any preliminary steps that are used to reduce the amount of information in the original data sets to a size more manageable for display to the human observer (for instance selecting out a single variable, functions of two variables, or higher level process such as feature space statistical pattern recognition). This paper describes a data visualization technique, named *isoluminance*, that encodes two variable data sets in in a single color image so that each variable and their combined spatial interaction can be understood by the visual mechanisms of hue and luminance perception, when displayed on computer workstations.

Initially, background work where color has been used for single and multivariable display, will be described. Also a distinction between the two ways variables are presented, discrete and continuous, is discussed. Then the design issues involved in generating a method that maps two variables to the perceived visual sensations of hue and luminance are described. An accurate (CIE human observer based) and approximate (heuristic) technique are described. Next, an implementation of the heuristic method for a standard workstation is described. Finally, we evaluate the technique after some preliminary trials using example medical images.

0-8194-1129-9/93/$6.00

2. BACKGROUND

2.1 Current methods for visualizing multivariable data sets

The method commonly used for displaying multivariable data sets is the side by side display of the two separate images. Other techniques have been used, including forms of temporal alternation (i.e. fading between or alternating between multiple variables) and relief maps, where one variable is shown in height, while the other variable is depicted as the surface intensity or color. Neither of these is in common use in radiology, although, cine, a form of alternation, has been used in some areas of medical imaging, especially for time sequences and occasionally for different MRI echoes.

2.2 Use of color in single variable display

Color has been used previously in attempts to improve human perception of the information in the image. Color has been applied to the single variable display problem in hopes of increasing the perceived dynamic range of color display devices. Several researchers have shown improved results using heated object (HOBS),[3] magenta,[3] and LOCS[4] pseudocolor mappings. Unfortunately, many early experiences of radiologists using color in medical image display were based on systems designed instead with ad hoc choices for pseudocolor maps resulting in a lack of acceptance of color, and in many cases a bias against the use of color in displays in general.[3,4]

2.3 Use of color in multivariable display

In radiology, there are examples of using color to show additional variables. One example is the overlaying of color on ultrasound images to show flow. In this case, however, the anatomical information under the colored areas is no longer visible. Work attempting to encode two variables into hue and luminance information has been done with nuclear medicine[5] and MR images.[6,7] These efforts, however, modeled only the mapping from data space to the hardware of the display system, leaving out the important consideration of human visual perception. Additionally none of the techniques provided for interactive user control of the mapping parameters (i.e. the downsampling from the variable spaces to the smaller number of levels in luminance or hue values).

Work in other fields, especially cartography and statistics, have dealt with the same problem of displaying multivariable data. In these fields, the problem of displaying *bivariate* maps (two variables at each map location) has been receiving attention since the Census Bureau produced and widely distributed some bivariate maps in the mid 1970s.[8,9,10] The visualization techniques in this area have not been based on encoding the two variables into hue and luminance, respectively. This is probably due to the differing requirements (the data sets usually have 3-6 levels per variable versus the 4096 levels common in medical images). Thus, work in this area has concentrated on trying to optimize the color choices for each of the possible combinations of one variable crossed with the other variable.

Finally, attempts have been made to encode three variables using hue, luminance and saturation.[11]

2.4 Continuous versus discrete depictions of variable spaces

The analysis of a variable or set of variables usually can be described as discrete or continuous. By discrete, we mean the variable space is segregated into separate identifiable groups. In the continuous case, the variable space is represented by smoothly continuous steps in a perceptual dimension. An example of a discrete representation from statistics might be to highlight areas where people made greater than $100,000 and paid no taxes. This resulting visualization would have four separate categories since each of the two variable spaces has been reduced to only two possible values. This type of example is often used to communicate a specific result or idea. On the other hand, the exploration, or less directed analysis of a data set, requires access to the full continuous variable spaces. Thus initially the expert user may examine the continuous data set, and upon arriving at a conclusion regarding the data set, condense it via a quantization, or data reduction step, to a form that more conveniently conveys the result to a specific audience. In the past, data visualization techniques have generally supported one or the other of these classes. With today's computers and high resolution monitors, it is now possible to design techniques that support both methodologies equally well at the same time.[8]

3. DESIGN

3.1 Driving Applications

Initial applications of isoluminance have been to medical image data sets that represent both anatomical and metabolic information, or structure and function, respectively. In magnetic resonance imaging we are investigating using MRS to detect phosphate or hydrogen nuclei concentrations to infer metabolic activity, and correlating this with standard T2 weighted images for the anatomical information.[12] A second example is the combination of emission data from a tumor imaging radiopharmaceutical with transmission data of the same nuclear medicine scan, as well as possibly with associated CT

images.[13,14] A third example is correlating lung function, depicted by nuclear medicine SPECT scan, with lung anatomy as seen on chest CT scans. In addition to the anatomical versus metabolic data sets, several other data sets were included in the analysis to test the applicability of isoluminance to other problems. Data sets in this group included standard T1/T2 MRI and census demographics for the USA. The MR anatomical image combined with MRS metabolic image were our primary example in development and testing of isoluminance.

3.2 Model for mapping from data space to perceptual space

Perceptually linearizing the mapping, or *linearization*, from the data space of the variable to the human observer's visual perceptual sensory space maximizes the information transfer to the human observer.[15,16,17,18] This simply means that to the human observer, equal proportional changes in the input variable should result in equal proportional changes in the visual sensation they perceive. Thus to most effectively represent a two variable space to the human observer, there should be a perceptual linear mapping from the data space of each variable to a separate sensory space of the human observer. By separate, we mean that if variable A is mapped to visual stimulus X, and variable B is mapped to visual stimulus Y, then changes in A should affect only X and not Y, and similarly changes in B should affect Y and not X. Vision research indicates that the hue and luminance information is processed fairly independently by the human brain,[19,20] and that human observers can quickly make decisions based on hue and luminance information presented simultaneously in a single image.[21,22] Thus the choice of encoding one variable as hue and the other as luminance.

The problem of mapping a two variable data set to the perceived channels of hue and luminance in a human observer is depicted in figure 1. Values for the number of levels in figure 1 are based on 12 bits of image contrast in medical images ($2^{12} = 4096$); 8 bits ($2^8 = 256$) for simultaneously displayable colors in the display system; the number of perceived luminance just noticeable differences, *JNDs* ; and the number of different hues on standard monitors. A human observer can perceive approximately 60-90 grey levels using a linearized grey scale on a standard workstation monitor.[23,4] The color gamut of a color CRT is very large, exceeding those of print and photography.[24] It is constrained by the number of different levels at which the RGB guns on the CRT can be driven. This is commonly 256 levels for each of R, G, and B, resulting in 2^{24} levels. While this overestimates the actual number of JNDs in hue as not all the produced hues could be distinguished by a human observer, the number $2^{24}/60$ provides a reasonable upper bound on the number of distinct hues that could be produced at a single luminance level.

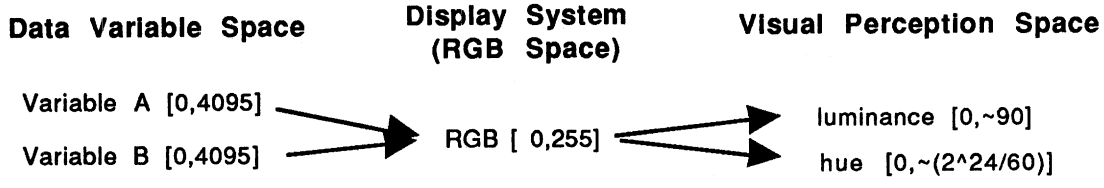

Figure 1 Data, RGB, and Perceptual spaces involved in mappings

As seen in figure 1, the overall mapping from the data space into the visual sensory space is actually two mappings because we are forced to work with the given characteristics of a specific monitor and video memory, or *display system*, on the workstation. This means the information in the data set will be quantized, at a possibly coarser representation, in RGB space. Additionally, the entire visual stimulus space for hue and luminance will not be available, only the portion that can be reproduced by the color monitor.

3.2.1 RGB space to perceptual space

The mapping from RGB space to perception space is constant because of the fixed set of visual stimuli the display system can reproduce. Thus with information on the specifics of the monitor characteristics and data on human observer hue and luminance perception, this relation can be defined statically for a given monitor system. While the mapping may be affected by room lighting conditions, etc.,[19] we considered these factors to be of secondary importance.

Thus, the choice in defining the structure of information in the RGB space and the mapping of data variable values to the RGB space determines the effectiveness of the technique. Appropriate choices should optimize the information transfer to the user, while inappropriate choices will reduce the information content as well as make the presentation more difficult to comprehend.

3.2.2 RGB structure

The best choice here is to completely define a mapping between RGB values and a perceptually linear color space. This would allow calculations of mapping from data variable space directly to perceptual space, with a final step of backprojecting the perceptual space values to RGB space for rendering on the specific display system. As a second choice, the RGB values could be represented by other color models or constructs that attempt to approximate the correct perceptual responses.

3.2.3 Data set to RGB mapping

The output of the video monitor is controlled by values placed in the computer's video memory. The structure of this memory dictates the spatial resolution (lines and columns of pixels on the display). The color resolution of a color display system is determined by the number of RGB combinations that can be simultaneously driven on the system. Generally the *depth* or number of bits of the memory element of the video memory determines the color resolution. For instance, most low end workstations have a single 8 bit byte as the memory representation, allowing 2^8 different color choices. In order to partition this space to encode the hue and luminance variables, the available color choices must be formatted to handle any hue value in combination with any luminance value. Thus, with an 8 bit system, for instance, one would have to choose the values h and l, where

h is the number of levels in the hue scale,
l is the number of levels in the luminance scale,
such that, $h*l <= 256$

to partition the 256 choices. Choosing h and l should be based on the characteristics of the three spaces (data, RGB, perceptual) depicted in figure 1. The task is to map the possibly larger number of levels in the two data space variables down to the number of levels of luminance and hue scales, respectively, in the RGB space. The number of levels in the perceptual space is important because, for instance, realizing that the observer may not perceive more than approximately 60-90 levels of luminance, it does not make sense to use luminance scales with more than 90 levels. In medical imaging the resolution of the input data is usually 256 to 4096 levels (8 to 12 bits of depth), although not all of this may be significant. Thus, high end workstations with 24 bit color display systems could support 2^{24} levels, which could be structured as 2^{12} luminance levels and 2^{12} hue levels allowing a one to one mapping from the two data variables ($2^{12} * 2^{12}$). On workstations with fewer levels, a compromise may be required. Given that only 60-90 levels of luminance are perceived, an initial choice might be to constrain the luminance levels to approximately 80, and determine if sufficient levels were left for depicting the variable represented by hue. On a low end 8 bit workstation further compromise in luminance levels would likely be required because even the choice of 80 levels for luminance leaves only 3 levels for hue.

3.3 User interface design for isoluminance model

Several important considerations are determined by the contents of the data sets, the capabilities of the display system and the preferences and abilities of the user. In many cases some information may be lost or reduced in content depending on these choices. Allowing interactive control over the parameters controlling these choices gives the user the ability to best choose those tradeoffs, as well as to use multiple visualizations in order to fully appreciate the data sets. The following attributes are the configurable parameters of the isoluminance model.

Which variable is mapped to hue and which to luminance, respectively

For medical image presentation with anatomical and metabolic information, we traditionally choose to display luminance as anatomy because this matches the radiologist's standard practice of viewing grey scale images where anatomy is depicted via the luminance of the grey scale.

Ability to fade between luminance depiction and hue depiction

In some cases the isoluminance technique may not be effective or the user may require assistance in understanding the resulting image. By providing the ability to fade between the two variables under user control, the observer can easily study either image independently, as well as slowly fade one of the variables in or out to assist in mentally comprehending the combined image.

Number of levels in luminance scale

For metabolic/anatomical combinations we attempted to maintain the number of levels as close as possible to what the user is capable of seeing, while not significantly degrading the metabolic information.

Number of levels in hue scale
In most metabolic/anatomical combinations the metabolic information could be represented with a small number of levels and still convey the appropriate information. Thus we usually chose to more significantly reduce the information content of metabolic activity compared to luminance information when required.

Choice of hue range
Because users prefer different hues, as well as perceive different hues better than others (especially color defective observers, such a dichromats, e.g. red green colorblind observers), it is very important to allow user choice of the hue range.

Choice of luminance range
While the choice of luminance range might also be considered for interactive control, we intentionally excluded it because a reasonable determination can be made from the display system characteristics combined with the hue scale choice information. That is, one generally wants as large a luminance scale as can be reproduced on the monitor given the chosen hue range.

Discrete versus continuous representations
Continuous representations of the variable are produced automatically because of the linearized mapping between the data variable and the perceptual space. Thus the important consideration is how to provide discrete representations while utilizing the same model and user interface. If the user is provided with the ability to identify any contiguous subset, or segment, of the input space, and to assign any hue range to this segment, then they can cause any segment to stand out visually via an appropriate hue range for that segment. This technique is referred to as labeling or highlighting.

4. IMPLEMENTATION

The initial development of the isoluminance model was done on a standard Unix workstation (Sun Sparc2) supporting the X Window System. This machine provides an 8 bit video buffer (256 levels in RGB space). While it may be more desirable to consider 24 bit workstations, these are not as ubiquitous as personal computers or workstations providing 8 bit support. Having only 256 levels to work with in RGB space meant compromising both the anatomical and metabolic variables. We felt that using less than 8 levels would would significantly compromise the metabolic information, so we chose to use 32 levels for the anatomical information and 8 for the metabolic information. While 8 levels of metabolic information is clearly inadequate in some cases, we felt it to be adequate for the metabolic cases we were studying. Choosing only 32 levels for anatomical detail represents a degradation of the information below the original acquisition (256 levels), and even below what the human observer can distinguish on the monitor (60-90), so it is likely that some fine detail would be lost.

Deciding upon 8 levels of hue to represent metabolic function and 32 levels of luminance to represent anatomical function dictates a structure like in fig 2. Figure 2 shows the 32 levels in horizontal rows, each having the same luminance values across the 8 vertical columns, which represent the 8 hue levels. The next step is producing a method for translating this 32x8 structure to perceived luminance and hue scales.

Optimal model
The best way would be to generate a mapping between a perceptually linear color space, for instance CIELUV,[25] and the RGB space of the monitor.[26,27,28] Then choose luminance and hue ranges and sample them at the desired resolution, for instance, 32 and 8, respectively, in CIELUV space. Using the CIELUV to RGB mapping for the display system, one can generate the 32x8 range of luminance and hue levels in RGB space. Then a linear mapping from the input data space to these RGB values would yield a perceptually linear mapping from the input data space to the perceptual space. A colorimeter or accurate specifications of the characteristics of the display system would be required to generate the perceptual to RGB space mapping.

Heuristic model
We chose to initially develop a simpler model, which would allow us to develop and test several isoluminance mappings. The heuristic was designed to allow easy construction of several isoluminance mappings based on interactive feedback from the observer as its only input.

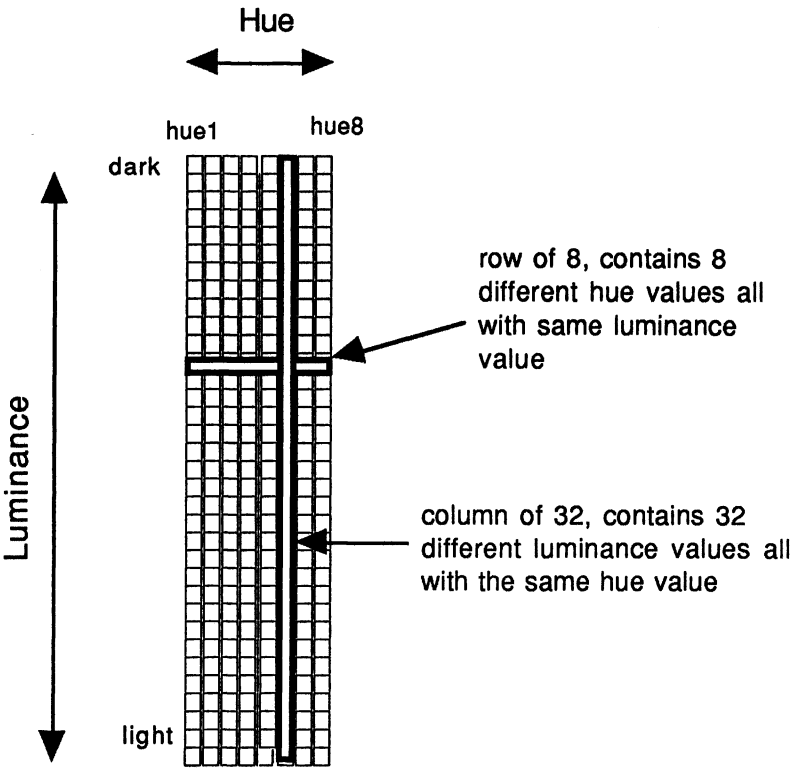

Figure 2 RGB structure for 32x8 isoluminance table

The same 32x8 table, as seen in figure 2, was utilized. This table is displayed to the observer on the screen. The hue range is designed to range from one primary to another primary. For example, consider red and green. Two curves are defined. One represents the percent contribution of the red gun at each of the 8 hues; the other curve represents the contributions of the green gun at the same 8 hues. Thus the two curves define 8 pairs of percent contributions from the combined red and green guns. These pairs define the hues at each of the 8 hue values for every luminance choice. The luminance range was simply chosen to range from the minimum digital driving value, 0, to the maximum digital driving value, 255. The luminance range increments in steps of 4 (i.e. the digital driving range, 256, divided by number levels in luminance, 32). Interactive slider bars were used to control each of the 8 control points on each of the primary hue contribution curves. Quick and convenient modification of the slider values allowed observers to choose the two hue curves so that the resulting 32x8 maps appeared perceptually to have the desired properties. Isoluminance maps that were combinations of two of the three primaries (R,G,B) could be easily generated with monotonic functions. Other hue ranges could also be generated, but less intuitively. Figure 3, for example, shows the red and green curves for the RG map. This map was generated by first choosing the endpoints of the hue range (most green and most red). Because the green phosphor on most monitors is brighter than the red phosphor when both are driven at their maximum digital driving level, we chose the ends of the hue scale to be

red end of hue scale = 100% red + S% green
green end of hue scale = T% green + 0% red
such that luminance(red end) = luminance (green end)

We then chose the percentage S to be small enough so that the hue was mainly red, but with enough green contribution to increase the brightness of the value since full red by itself is significantly dimmer than full green. In figure 3, for instance, S is 22% and T is 88%.

Once the endpoints of the range were fixed we would display the entire 32x8 table and interactively adjust the sliders to achieve a satisfactory selection of 8 hue steps across each row, while maintaining similar luminance values for each member of the row.

The choice of 8 control points defining the curve also provided a mechanism for highlighting. By significantly changing one control point, the segment represented by that control point could be made to stand out visually.

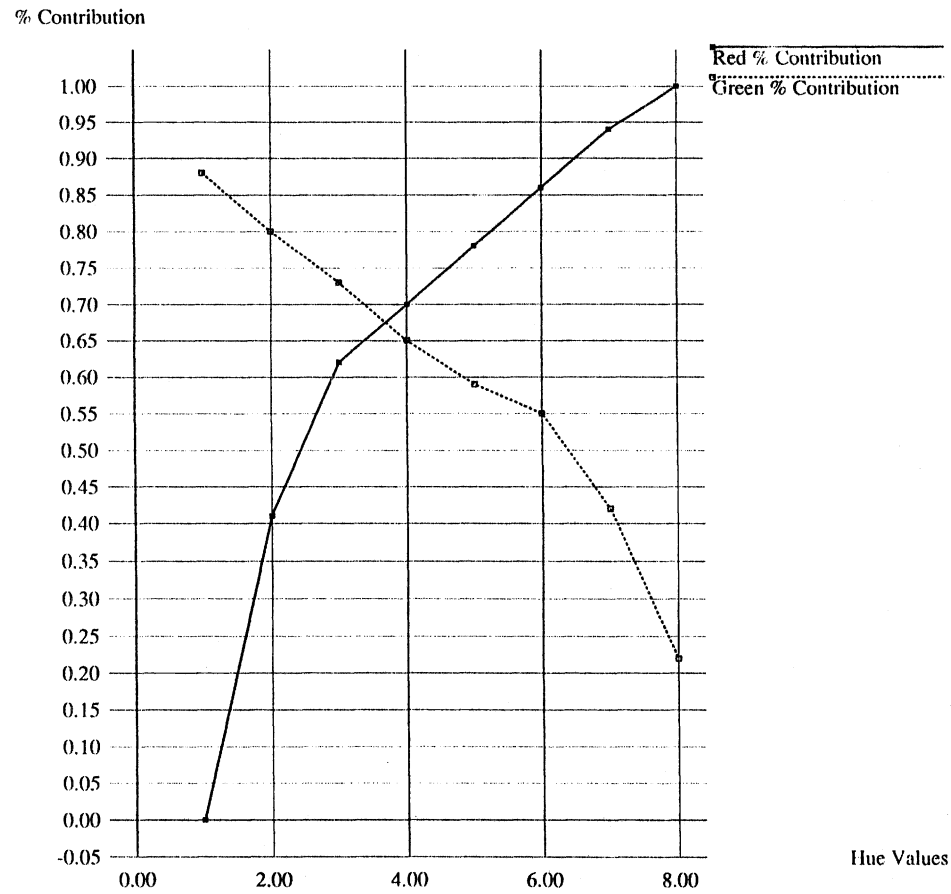

Figure 3 Hue Control Curves for Red-Green isoluminance map

Another interaction easily supported was the fade in and fade out of variables. We wanted to be able to show just the luminance variable or just the hue variable via an interactive slider control. Fading towards showing luminance was accomplished by simply flattening the curves toward their combined mean. Fading towards hue was accomplished by simply shrinking the luminance range. We chose to shrink the luminance not to a center mean luminance, but instead to about the 75th percentile between the bottom and top values to best visualize the hues. This is because the hue differences tend to be perceived better at higher luminances, except for luminances at the top end of the monitor where the range of displayable hues diminishes.[28,26]

5. EXPERIENCE

Seven observers, including radiologists, computer scientists and graduate students, used the system. They were presented with the original data presented as two grey scale images and the combined color image at the same time. We asked them to compare the presentations and to specify whether they felt they could comprehend the information from each of the two variables when they were presented as hue and luminance. Additionally we asked them to experiment with the choice of the isoluminance mapping (by modifying the hue curve sliders), to fade from one variable to the other, to try highlighting segments, and to reverse the initial choice of which variable mapped to hue and vice-versa. From this we learned several things regarding the effectiveness of isoluminance in general and also with regard to the specifics of our heuristic model. These are summarized below.

Downsampling to smaller RGB space

We found the choice of the number of levels in the hue and luminance scales to be very dependent on the data set and the visual task. For the anatomical and metabolic combinations in MR/MRS, we found that 256 levels for the hue and luminance scales in RGB space was often sufficient, and that the choice of 32 levels for the luminance scale and 8 levels for the hue scale was the most preferred. The users indicated a higher level of satisfaction with the mapping if we allowed them control of the downsampling rather than fixing the mappings from the data space to the RGB space. Clearly, however, it would be more ideal to have access to enough levels in the RGB space to not have to downsample. We compared the 8 bit display system with a 24 bit display system(12 bits for luminance and 12 bits for hue) and found that most observers were able to detect differences in the luminance scale and that all observers were able to see differences in the hue scale. The latter is expected because in the 8 bit display system we sampled the one variable to only 8 levels of hue, as compared with 4096 levels of hue when displayed on the 24 bit display system.

Variable encoded as hue

In situations, like the metabolic/anatomic combination, where the metabolic variable tended to be simple in structure, the hue presentation was effective at conveying the associated variable and its position relative to the anatomical information. Most observers had more difficulty with other combinations and often the variable represented by hue was not seen as accurately, or as easily as in the separate grey scale image. Several common factors were present in cases where users found the hue presentation did not convey the information as well as a separate display of the image. First, and most common, was when the variable represented by luminance had many different levels, and contained significant amounts of high frequency information (large changes in value between nearby pixels). Additionally, if the hue variable also had significant amounts of high frequency information this worsened the problem. Also, when the variable was similar to the luminance variable in structure, observers often had difficulty sorting out the two sets of information, again, especially in the presence of high frequency information. A common example of this was MR images with both showing anatomy, for instance, T1/T2 image pairs. Finally, hue differences at low luminances were not perceived well by the observers. This resulted in our choosing to shift the luminance scale to start not at the darkest luminance of the monitor but instead at luminance values where most observers could more easily distinguish hue differences. Understanding the hue variable was a significant problem overall, in that observers often felt that while they could study the color image and appreciate the variable encoded as hue, that many times this information did not "stand out," or would require effort to comprehend.

Variable encoded as luminance

Observers felt the presentation of the variable represented by luminance was very good. The major difference cited by observers was when the variable space was downsampled significantly for presentation in luminance scales with a limited number of levels. We were especially sensitive to this problem because one of our initial goals was to maintain the mapping of anatomy as luminance in the anatomical/metabolic studies to capitalize on the experience of observers in reading grey scale (i.e. luminance encoded) images. The luminance scales generally used, however, differed in two ways. The first difference was the shifting upwards of the bottom of the luminance range due to the decreased appreciation of darker hues. The second difference was that the bright end of the luminance scale would have only a single gun driven to its full capacity, as opposed to all three guns in a grey scale luminance map.

While the first problem was noted by observers initially, it was not noted as significant after we shifted the low end of the luminance scale. We also used low ambient light conditions when viewing the images which likely helped as well. Observers only noted this as a problem, after the change, when using the RB (dimmest) mappings. The second problem, combined with the shifting up of the low end of the luminance scale, resulted in only a portion of the monitor's luminance range being utilized. Because the human observer can only recognize 60-90 luminance levels for a given display system, if we were cropping this range to even fewer levels we might be degrading the representation of the variable mapped to the luminance scale. For the images in our sample, however, the users felt the variable mapped to the luminance scale was well represented.

Which variable encoded as luminance versus hue

For medical images depicting anatomical structure versus metabolic function observers always preferred anatomy to be represented by luminance. This is most likely because all observers had previous experience seeing grey scale images where anatomy is shown via luminance scales. In other situations it strongly depended on the data sets involved.

Fading

The ability to fade between the images proved to be a significant tool in cases where the initial presentation was not easily comprehended. The users would often fade back and forth, or start with one variable only and slowly add the second variable, apparently to assist in integrating the variables in their mind by slowly integrating them on the screen.

Highlighting

In addition to user control of fading, the ability to highlight areas was deemed extremely useful. This appeared to be a good way to perform discrete labeling or characterization of one of the data sets. Because the highlighting was controlled by changing the definition of a control point on the hue control curve in an arbitrary fashion, such changes resulted in a loss of the isoluminance properties. A better methodology that would allow highlighting while maintaining the isoluminance mapping, for instance by highlighting with a hue outside of the current hue scale range, was deemed to be an improvement. Additionally, a mechanism for choosing segments of the variable space for the highlighting was mentioned as a desired function. This would entail providing support for a user selected segment instead of only allowing the user to change one of the 8 predefined segments, in the instance of the 32x8 example.

Heuristic versus best model

Probably the biggest shortcoming of the initial system was the heuristic used to generate the hue and luminance mappings. While the use of the heuristic allowed the development and testing of a technique having approximately isoluminance properties with the ability for interactive control of parameters of the technique, the heuristic had several significant shortcomings:

1) chosen mappings were only rough approximations to proper isoluminance maps. The observers indicated that while the hue columns seemed consistent in hue, the luminance values across rows did not seem to be identical. Follow-up measurements of the 8 hues of a few isoluminance rows often showed significant differences in luminance values for members of the same row when measured with a photometer. Figure 4 shows an example of differences in luminance values from one row of the RG map. Since the luminance levels do not directly represent the perceived visual responses, one observer was additionally tested using a standard psychophysical luminance two patch matching task.[29] Results correlated with the photometer measured differences, indicating differences in luminances between two members from the same row, and usually showing several JND luminance steps between the two.

2.7	2.1	1.8	1.5	1.4	1.4	1.2	1.3

Table 1 Photometer measured values (in footlamberts) for 8 elements of row 8 of 32x8 RG isoluminance table.

2) The selection of hue ranges was limited. The choice was limited to combinations of two guns, making choices besides RG, GB, and RB more difficult to accomplish. Additionally, many possible isoluminant mappings could not be defined by this method because it is based on the combination of only two of the RGB guns.
3) The users often indicated a desire to define their own mapping, but were frustrated when attempting to do so with the existing method. A special case was users who were RG colorblind, as the RG mapping was usually the preferred mapping. RG colorblind observers tended to be more dissatisfied with the mappings, indicating a desire to choose their own mapping. This suggests an important attribute is to provide interactive control of the hue range, or at least a larger selection of good (prechosen) mappings.
4) The problem with the reduced luminance range described above in this section would disappear if all three guns could be utilized, which was not possible with the heuristic.

6. CONCLUSION

We had positive responses using isoluminance for multivariable problems with one data set representing anatomical function and the other data set representing metabolic function. We had less success with data sets that convey similar information or where both data sets were mainly comprised of high frequency information. In the latter case, users relied on the ability to fade back and forth between the luminance and hue representations indicating that fading was probably as effective as the inherent advantages of the combined hue and luminance depictions. While the heuristic color model worked reasonably well for a limited

set of isoluminant mappings, it is clear that using an accurate perceptual to RGB mapping would improve on the quality of the isoluminance mappings as well as possible choices of mappings.

7. FUTURE WORK

We are in the process of obtaining a colorimeter to generate RGB to CIELUV mappings for our monitors. The next iteration of isoluminance will incorporate this CIELUV to RGB mapping instead of the heuristic used in the first iteration. Additionally the user interface is being redesigned to allow arbitrary segment definition as well as *hue outside of hue scale* highlighting so that highlighting will work while maintaining the isoluminance properties. Finally, interactions for all the parameters described in section **3.3** will be redone based on the new CIELUV to RGB mapping. Finally, we expect to carry out a controlled observer experiment to evaluate the effectiveness of isoluminance versus side by side display for MR/MRS image sets.

Special Note
No color pictures were included in this paper because of the significant cost involved. If you would like to see images generated with the isoluminance technique, please contact the author directly at the address on the title page or via electronic mail at *bmh@rad.unc.edu.*

8. ACKNOWLEDGEMENTS

This work was supported in part by NIH grant R01 CA 44060.

9. REFERENCES

1. McCormick, BH, DeFanti, TA, Brown, MD, Visualization in Scientific Computing, *Computer Graphics 21, 6* (November 1987), ACM SIGGRAPH: New York, pp 3.

2. McCormick, BH, DeFanti, TA, Brown, MD, Visualization in Scientific Computing, *Computer Graphics 21, 6* (November 1987), ACM SIGGRAPH: New York, pp A1-A15.

3. Pizer, SM, Zimmerman, JB, Johnston, RE, "Contrast Transmission in Medical Image Display", *Proceedings 1st International Symposium Medical Imaging and Interpretation* (ISMIII 82), pp. 2-9, 1982.

4. Levkowitz, H, Herman GT, "Color Scales for Image Data", *IEEE Computer Graphics and Applications*, Vol 12, No 1, pp72-80, Jan 1992.

5. Crowe, EJ, Sharp, PF, Undrill, PE, Ross, RGB, "Effectiveness of colour in displaying radionuclide images", *Med. & Biol. Eng. & Computing*, 26, pp. 57-61, 1988.

6. Wells, MG, Sharp, PF, Law, ANR, "Principles and appraisal of combined images in NMR", *Med. & Biol. Eng. & Computing*, 27, pp. 277-280, 1989.

7. Weiss, KL, Stiving SO, Herderick, EE, Cornhill, JF, Chakeres DW, "Hybrid Color MR Imaging Display", *AJR*, 149, pp. 825-829, 1987.

8. Dunn, R, "A dynamic approach to two-variable color mapping", *The American Statistician*, Vol. 43, No. 4, pp. 245-252, November 1989.

9. Olson, JM, "Spectrally encoded two-variable maps", *Annals of the Association of American Geographers*, Vol. 71, No. 2, pp.259-276, June 1981.

10. Trumbo, BE, "Theory for coloring Bivariate statistical maps", *The American Statistician*, Vol. 35, No. 4, pp. 220-226, November 1981.

11. Levkowitz, H, Herman, GT, "Color in multidimensional multiparameter medical imaging", *COLOR research and applications*, 11(Supplement), pp. S15-20, 1986.

12. Weiner, MW, "Clinical magnetic resonance spectroscopic imaging", Field Strength, Vol. 2., No. 1, pp. 1-20, March, 1990.

13. Frey, EC, Tsui, BMW, Perry, JR, "Simultaneous Acquisition of Emission and Transmission Data for Improved Thallium-201 Cardiac SPECT Imaging Using a Technetium-99m Transmission Source," The Journal of Nuclear Medicine, Vol. 33, No. 12, pp. 2238-2245, December 1992.

14. Gilland, DR, Tsui, BMW, Metz, CE, Jaszczak, RJ, Perry, JR, "An Evaluation of Maximum Likelihood-Expectation Maximization Reconstruction for SPECT by ROC Analysis," The Journal of Nuclear Medicine, Vol. 33, pp. 451-457, March 1992.

15. Ji, T, Roehrig, H, Blume, H, Guillen, J, "Optimizing the display function of display devices", *Proceedings Medical Imaging VI: Image Capture, Formatting*, and Display, Vol 1653, pp. 126-139, February 1992.

16. Briggs, SJ, "Photometric technique for deriving a "best gamma" for displays", *Advances in Display Technology*, SPIE 199, pp 134-145, 1979.

17. Pizer, SM, "Intensity mapping to linearize display devices", *Computer Graphics and Image Processing*, 17, pp 262-268, 1981.

18. Johnston, RE, Zimmerman, JB, Rogers, DC, et al., "Perceptual Standardization", *PACS III for Medical Application*, SPIE 536, pp 44-49, 1985.

19. Levkowitz, H, Robertson, PK, Rogowitz, BE, "Color Theory and Models for Computer Graphics Visualization", Course 10, SIGGRAPH 1992, pp 18-22, 40-44, 1992.

20. Burns, B, Shepp, BE, "Dimensional interactions and the structure of psychological space: The representation of hue, saturation, and brightness", *Perception and Psychophysics*, 43 (5), pp 494-507, 1988.

21. D'Zmura, M, "Color in visual search", *Vision Research*, Vol. 31, No. 6, pp. 951-966, 1991.

22. Nagy, AL, Sanchez, RR, "Chromaticity and Luminance as Coding Dimensions in Visual Search", *Human Factors*, 34(5), pp 601-614, 1992.

23. Pizer, SM, Chan, FH, "Evaluation of the number of discernible levels produced by a display", *Information Processing in Medical Imaging*, R DiPaola and E. Kahn, Eds., Editions INSERM, Paris, pp. 561-580, 1980.

24. Tannas, LE, "Color in electronic displays", *Physics Today*, pp. 55, December 1992.

25. Commission Internationale de l'Eclairage (CIE). CIE recommendations on uniform color spaces--color difference equations, psychometric color terms. CIE Publication (15, (E-1.3.1) 1971/(TC-1.3) 1978, Supplement No. 2), pp. 9-12, 1978.

26. Wyszecki, G, Stiles, WS, Color Science: Concepts and Methods, Quantitative Data and Formulae, pp. 117-248, 2nd Edition, John Wiley and Sons, New York, 1982.

27. Robertson, PK,"Visualizing Color Gamuts: A user interface for the effective use of perceptual color spaces in data displays", *IEEE Computer Graphics and Applications*, Vol. 8, No. 5, pp. 50-64, September 1988.

28. Taylor, JM, Murch, GM, McManus, PA, "TekHVC: A uniform perceptual color system for display users", *Proceedings of the SID*, Vol. 30/1, pp. 15-21, 1989.

29. Boyton, RM, Human Color Vision, Holt Rinehart Winston, New York, 1979.