Automatic Perceptual Color Map Generation for Realistic Volume Visualization

Jonathan C. Silverstein, **Nigel M. Parsad**, and **Victor Tsirline**
Departments of Surgery, Radiology, and Computation Institute, University of Chicago and Argonne National Laboratory

Abstract

Advances in computed tomography imaging technology and inexpensive high performance computer graphics hardware are making high-resolution, full color (24-bit) volume visualizations commonplace. However, many of the color maps used in volume rendering provide questionable value in knowledge representation and are non-perceptual thus biasing data analysis or even obscuring information. These drawbacks, coupled with our need for realistic anatomical volume rendering for teaching and surgical planning, has motivated us to explore the auto-generation of color maps that combine natural colorization with the perceptual discriminating capacity of grayscale. As evidenced by the examples shown that have been created by the algorithm described, the merging of perceptually accurate and realistically colorized virtual anatomy appears to insightfully interpret and impartially enhance volume rendered patient data.

Keywords

Volume visualization; Perceptual color maps; Surgical planning; Radiological anatomy; Virtual Reality

1. Introduction

Computed tomography (CT) scanners measure the linear x-ray radiodensity attenuation values of body tissues in Housfield units (HU). Transformation of the resulting data into two-dimensional (2D) CT slice images is routine and ubiquitous. Historically, mapping HU to a range of grayscale values to render tissue density was the only color map visualization option for medical image analysis. An accordingly broad scope of diagnostic imaging tools and techniques based upon grayscale two-dimensional (2D) image interpretation was thus established. Grayscale offers excellent perceptual discrimination for multidimensional visualization because it is monotonically increasing in luminance across its entire spectrum.

For more than a decade, automatic volume rendering of realistic three-dimensional (3D) objects directly from 2D CT slice images has been technically feasible and envisioned for real-time clinician interaction and treatment planning [1]. Intricately detailed volume reconstructions of large CT slice image datasets are now possible with the advent of high-resolution, multi-

Corresponding Author: Jonathan C. Silverstein, Computation Institute, University of Chicago, Research Institutes, Suite 405, 5640 South Ellis Avenue, Chicago, Illinois 60637, Phone (773) 834-8200 Fax (773) 834-8400, E-mail: jcs@uchicago.edu.

detector CT scanners. Inexpensive, high performance graphics processing units, particularly in parallel rendering architectures, allow for interactive manipulation and 24-bit full color representation of these data intensive volume visualizations [2]. In our experience teaching anatomy and in surgical treatment planning using these technologies, we find grayscale, despite its perceptual advantages, to be a severe distraction from the realism made possible by modern volume rendering, especially in stereoscopic visualization systems [3,4].

Some diagnostic radiological workstations even offer full color options for volume rendering CT data. However, many alternative color maps have little additional diagnostic or intuitive value relative to grayscale. Certain color maps can actually prejudice the interpretation of 2D/3D data [5]. A common example is spectral hue colorization, which is also often seen representing temperature range on weather maps [6]. Other preset colorization algorithms are apparently aesthetic creations with no apparent theoretical or pragmatic basis for enhancing data visualization or analysis. These limitations with existing color maps, coupled with our need for realistic volume rendering for teaching and surgical planning, motivated us to explore whether it was possible to automatically generate color maps that would permit simultaneously anatomic realism and the perceptual discriminating capacity of grayscale.

2. Color Theory Primer

2.1 Color Perception

In order to understand the complexity of using perceptual rules to create a color map with anatomically realistic hues, a brief discussion deriving the relation of human physiology and perception to color theory is required. Human color vision is selectively sensitive to certain wavelengths over the entire visible light spectrum. Furthermore, our perception of differences in color intensity is non-linear between hues. As a result, color perception is a complex interaction incorporating the brain's interpretation of the eye's biochemical response to the observed spectral power distribution of visible light. The *spectral power distribution* is the incident light's *intensity* per unit wavelength and is denoted by $I(\lambda)$. $I(\lambda)$ is a primary factor in characterizing a light source's true brightness and is proportional to $E(\lambda)$, the energy per unit wavelength. The sensory limitations of retinal cone cells combined with our non-linear cognitive perception of $I(\lambda)$ are fundamental biases in our conceptualization of color.

2.2 Physiology of color vision

Cone response is described by the trichromatic theory of human color vision. The eye contains 3 types of photoreceptor cones that are respectively stimulated by the wavelength peaks in the red, green or blue bands of the visible light electromagnetic spectrum. As such, trichromacy produces a weighted sensitivity response to $I(\lambda)$ based upon the RGB primary colors. This weighting is the *luminous efficiency function* ($V(\lambda)$) of the human visual system. The CIELAB color space recognizes the effect of trichromacy on true brightness via its *luminance* (Y) component. Luminance is the integral of the $I(\lambda)$ distribution multiplied by the luminous efficiency function and can be summarized as an idealized human observer's optical response to the actual brightness intensity of light [7]:

$$Y = \int I(\lambda) V(\lambda) \text{ where: } \sim 400nm < \lambda < 700nm \tag{1}$$

2.3 Psychophysics and perceived brightness

Empirical findings from the field of visual psychophysics show that the human perceptual response to luminance follows a compressive power law curve [8]. As luminance is simply the true brightness of the source weighted by the luminosity efficiency function, it follows that we perceive true brightness in a non-linear manner. A percentage increase in the incident light

intensity is not cognitively interpreted as an equal percentage increase in perceived brightness. CIELAB incorporates this perceptual relationship in its *lightness* component L*, which is a measure of *perceived brightness* [9]:

$$L^* = 116 (Y/Y_n)^{1/3} - 16 \text{ where: } 8.856x10^{-2} < Y/Y_n \tag{2}$$

Here Y_n is the luminance of a white reference point in the CIELAB color space. The cube root in the luminance ratio Y/Y_n approximates the compressive power law curve, e.g., a source having 25% of the white reference luminance is perceived to be 57% as bright. Note that for very low luminance, where the relative luminance ratio is lower than $8.856x10^{-2}$, L* is approximately linear in Y. To summarize, Y and L* are sensory defined for the visual systems of living creatures and light sensitive devices whereas $I(\lambda)$ is an actual physical attribute of electromagnetic radiation.

3. Background

3.1 Advantages of Realistically Hued Color Maps

Color realism is essential to surgeons' perceptions as they predominantly examine tissue either with the naked eye or through a camera. A known visualization obstacle in the surgical theater is that in a bloody field, it is difficult to see even in the absence of active bleeding or oozing. A simple rinsing of the field with saline brings forth an astounding amount of detail. This is because dried blood covering the tissues scatters the incident light, obscures the texture, and *conceals color information*. All of the typical color gradients of yellow fat, dark red liver, beefy red muscle, white/bluish tint fascia, pale white nerves, reddish gray intestine, pink lung, and so on become a gradient of uniform red which is nearly impossible to discriminate with a naked eye. This *lack of natural color gradients* is precisely the reason why grayscale and spectral colorizations cannot provide the perceptive picture of the anatomy no matter how sophisticated the 3D reconstruction.

Realistic colorization is also useful in rapidly identifying organs and structures for spatial orientation. Surgeons use several means for this purpose: shape, location, texture, and color. Shape, location, and, to some extent texture, are provided by 3D visualization in grayscale. However, this information may not be sufficient in all circumstances. Specifically, when looking at a large organ in close proximity, color information becomes invaluable and realism becomes key. In order to ensure real colors are represented in video-endoscopic surgery, every laparoscopist begins the case by white balancing the camera inside the abdomen. A surgeon's visual perception is entrained to the familiar colors that remain relatively constant between patients. Surgeons do not consciously ask themselves what organ corresponds to what color.

However, when looking at a grayscale CT image, one has to look at the shade of the object relative to other known structures in order to identify it rather than its inherent hue (see Figure 1 and Figure 2). Structures with no inherent shape, location, or texture are particularly challenging. For example, free fluid in the peritoneal cavity may be bloody, purulent, or serous, each of which have different densities but all appear as a medium gray in a typical abdominal window. Radiologists often explicitly measure the density of peritoneal fluid in Hounsfield units via an image query tool since the shape, location and texture are not sufficiently informative and there is no color information in grayscale.

The advantages of realistically hued color maps appear to apply particularly to anatomy education and to three-dimensional stereoscopic visualization environments for surgical planning. In both circumstances, simulating the appearance of live tissue is deemed extremely

advantageous by the participants [personal experience]. These advantages may also, when feasible, apply to many other circumstances, including possibly in diagnostic radiology, and other non-medical domains.

3.2 Complexity of 2D/3D Anatomic Visualization

In a typical full body scan, the distribution of HU ranges from -1000 for air to 1000 for cortical bone [10]. Distilled water is zero on the Hounsfield scale. Since the density range for typical CT datasets spans approximately 2000 HU and the grayscale spectrum is limited by typical computer monitors to the 256 grays in the RGB color space, radiologists are immediately faced with dynamic color range challenge for CT image data. Mapping 2000 density values onto 256 shades of gray results in an under-constrained color map. Lung tissue, for example, cannot be examined concurrently with the cardiac muscle or vertebrae in grayscale because the thoracic density information is spread across too extreme a range. Radiologists use techniques such as density windowing and opacity ramping to interactively increase density resolution. However, just as it is impossible to examine a small structure as high zoom without losing the rest of the image off the screen, it is impossible to examine a narrow density window without making the surrounding density information invisible. Vertebral features are lost in the lung window and vice-versa.

This problem compounds itself as you scroll through the dataset slice images. A proper window setting in a chest slice may not be relevant in the colon, for example. One must continually window each dissimilar set of slices to optimize observation. Conventional density color maps can accomplish this but the visualizations are often unnatural and confusing – it is often easier to look at them one organ at a time in grayscale. Fortunately radiologists can focus their scanning protocols on small sections of the anatomy where they can rely on known imaging parameters to compress dynamic range and systematically review the results. However, *without radiology training and experience*, it is often obscure which view or parameters optimizes seeing specific anatomical features, let alone attempting to reconstruct mentally the entire relevant region of anatomy rather than look at a realistic representation.

This problem is exacerbated in 3D as the imaging complexity of the anatomical visualization substantially increases with the extra degree of spatial freedom. For example, a volume rendered heart contains tissue densities including fat, cardiac muscle, and cardiac vasculature in a relatively compact space. If the region of interest around the heart is extended, then vertebrae, bronchia, and liver parenchyma are immediately incorporated into the visualization. This increases the parameter space of visualization data considerably with just a small spatial addition of volume rendered anatomy. A modern computer display can produce millions of colors and thus overcome the challenge of wide dynamic range HU visualizations. This is especially important in 3D volume renderings because large anatomical regions are typically viewed. Intuitive mapping of color to voxel density data becomes a necessity, because simply mapping an ad hoc color map onto human tissue HU values does not guarantee insightful results.

Anatomically realistic color maps also allow for a wider perceivable dynamic visualization range at extreme density values. Consider that volume reconstruction of the vertebrae, cardiac structure, and airfilled lung tissues can be displayed concurrently in fine detail with realistic colorization, i.e., the thoracic cardiac region would find the vertebrae mapped to white, cardiac muscle to red, fat to yellow, and the lung parenchyma to pink with air transparent (see Figure 1 and Figure 2).

Object discrimination on the basis of color becomes especially important when clipping planes and density windowing are used to look at vessels within solid organs or to "see through" the thoracic cage, for instance. These techniques allow unique visualizations of intraparenchymal

lesions while mentally mapping the vessels and ducts at oblique angles. However, one can easily lose orientation and thus the *spatial* relationships between structures in such views. Color realism of structures maintains this orientation for the observer as natural colors obviate the need to question whether the structure is a bronchus or a pulmonary artery or vein.

3.3 Advantages of Luminance Controlled Color Maps

Despite the advantage of realistic colorization in the representation and display of volume rendered CT data, grayscale remains inherently superior with regard to two important visualization criteria. Research in the psychophysics of color vision suggests that color maps with *monotonically increasing* luminance, such as grayscale, are perceived by observers to naturally enhance the spatial acuity and overall shape recognition of facial features [11]. Although these studies were only done on two-dimensional images, due to the complexity of facial geometry, we feel this study is a good proxy for pattern recognition of complex organic shapes in 3D anatomy.

Findings also suggest that color maps with monotonically increasing luminance are ideal for representing interval data [6]. The HU density distribution and the Fahrenheit temperature scale are examples of such data. Interval data can be defined as data whose characteristic value changes equally with each step, e.g., doubling the Fahrenheit temperature results in a temperature twice as warm [12]. A voxel with double the HU density value relative to another is perceptually represented by a color with proportionally higher luminance. It follows from Eq. (2) that the denser voxel will also have a higher perceived brightness. Grayscale color maps in medical imaging and volume rendering are therefore perceptual because luminance, and thus perceived brightness, increases monotonically with tissue density pixel/voxel values.

Secondly, whether the HU data spans the entire CT range of densities or just a small subset (i.e. "window") of the HU range, the range of grayscale's perceived brightness (L*) is *maximally* spread from black to white. Color vision experiments with human observers show that color maps with monotonically increasing luminance and a maximum perceived brightness contrast difference greater than 20% produced data colorizations deemed most natural [11]. Color scales with perceived brightness contrast below 20% are deemed confusing or unnatural regardless if their luminance increases with monotonically with the underlying interval data. From these two empirical findings, it appears that grayscale colorization would be the most effective color scale for HU density data. However, for anatomical volume rendering, despite being comfortable to view, grayscale doesn't convey realism, thus leading to a distracting degree of artificialness in an otherwise immersive visualization.

4. Method

4.1 Stereoscopic Volume Visualization Engine and Infrastructure

We have built a highly flexible environment in which to interactively experiment with stereoscopic volume rendering and color maps [3]. High-resolution volume datasets are loaded without preprocessing by a parallel graphics processing unit visualization system. The system is open architecture, open source, utilizes commoditized hardware, and available for others to use. It is integrated with Access Grid (www.accessgrid.org) for multi-location, multi-user collaborative rendering and control. The distributed, collaborative functionality underlying the visualization architecture necessitated a user interface design principle focused on simplicity to drive powerful functions underneath.

The software runs on a nine-node, high performance graphics computing cluster. Each node runs an open source Linux OS and is powered by an AMD Athlon 64 Dual Core 4600+ processor. The volume rendering duties are distributed among eight "slave" nodes. A partial

3D volume reconstruction of the CT dataset is done on each slave node by an Nvidia 7800GT video gaming card utilizing OpenGL/OpenGL Shader Language. The remaining "master" node assembles the renderings and monitors changes in the rendering's state information. Each eye perspective is reconstructed exclusively among half of the slave nodes, i.e., four nodes render the left or right eye vantage point respectively. The difference in each rendering is an inter-ocular virtual camera offset that simulates binocular stereovision. Both eye perspectives are individually outputted from the master node's dual-head video card to their own respective video projector. The projectors overlap both renderings on a 6'x5' projection screen. Passive stereoscopic volume visualization is very inexpensively achieved when the overlapped renderings are projected through appropriately polarized filters and viewed through matching polarized glasses. The virtual environment is controlled via a front-end graphical user interface or GUI (Figure 3). The volume and GUI are controlled via a single button mouse. The GUI's available features include those typical of volume rendering on radiological workstations: such as rotation, zooming, panning, multi-plane clipping, HU density windowing, and color map selection.

4.2 Volume Rendering and Automated Colorization of Hounsfield CT Data

Full body CT scans produce HU density distributions with characteristic density peaks that correspond to specific anatomic features. Such a distribution is shown in Figure 4. The blue line, which is superimposed over our generic, realistic color map described in Table 1, Section 5.1, is the HU distribution extracted from a nearly full body CT scan. Classical teaching in radiology informs us that four regions are delineated (air, fat, tissue, and bone). These peaks can be seen in Figure 4.

HU distributions are similar in data representation to scalar temperature fields used for national weather maps. Both datasets provide scalar data values at specific locations. The spatial resolution of a temperature field is dependent on the number of weather stations you have per square mile. The spatial resolution of the HU distribution depends on the resolution of the CT scanner. State of the art 64 detector scanners can determine HU values per axial image slice for areas less than a square millimeter. During volume rendering, the distance between CT axial image slices determines the Z-axis resolution. The resulting 3D voxel inherits its interpolated HU value from corresponding 2D slices. Depending on rendering algorithms used, the HU voxel value can be continuously changing based on the gradient difference between adjacent slice pixels. The shape of the HU distribution is dependent on what part of the body is scanned much like the shape of a temperature distribution depends on what area of the country you measure.

RGB, the familiar red-green-blue additive color space utilizes the trichromacy blending of these primary colors in human vision. Thus this color model can be represented as a 3 dimensional vector in color space with each axis represented by one of the RGB primary colors and with magnitudes ranging from 0 to 255 for each component. In general, in a typical full color volume rendering, a look-up-table (LUT - such as Table 1) relates RGB values to HU value intervals. The varying opacity window (described below) then typically determines the specific color's HU correspondence by affecting the color's transparency and also the specific HU range over which the LUT is spread. In addition, in renderings using the perceptual luminance-matching method which is the focus of this paper, each RGB value in the look-up-table itself, not just its opacity or HU correspondence, is also adjusted to match the luminance prescribed by the transfer function for that point of the selected window were it to have been grayscale.

4.3 Opacity Windowing

RGBA adds an alpha channel component to the RGB color space. The alpha channel controls the transparency information of color and ranges between 0% (fully transparent) to 100% (fully opaque). Our color algorithm is integrated with several standard opacity ramps that modify the alpha channel as a function of density for the particular window width displayed. Opacity windowing is a necessary auto-segmentation tool in 2D medical imaging. We have extended it to volume rendering by manipulating the opacity of a voxel as opposed to a pixel. For example, the abdomen window is a standard radiological diagnostic imaging setting for the analysis of 2D CT grayscale images. The abdomen window spans from -135HU to 215HU and clearly reveals a wide range of thoracic features. The abdomen window linear opacity ramp renders dataset voxels valued at -135HU completely transparent, or 0% opaque, and the voxels valued at 215HU fully visible, or 100% opaque. Since the ramp is linear, the voxel at 40HU is 50% transparent. All other alpha values within the abdomen window would be similarly interpolated. Voxels with HU values outside of the abdomen window (below 135HU or above 215HU) would have an alpha channel value of zero, effectively rendering then invisible.

Although the linear opacity ramp is the only transfer function whose characteristics are explored in this colorization effort, we have optionally employed several typical non-linear opacity functions to modify the voxel transparency including Gaussian and logarithmic ramps; the colorization algorithms discussed can be usefully applied with those transfer functions as well. Similarly we expect that other transfer functions we have not employed, such as those of derivatives of HU, may also be usefully combined with the perceptual luminance-matching algorithm.

5. Results

5.1 Anatomically Realistic Base Color Map

In order to produce a density based color map scheme that would mimic natural color, we first assigned a natural color range for each of the primary density structures of the human body and then interpolated between them. Selection of realistic, anatomical colors was done by color-picking representative tissue image data from various sources such as the Visible Human Project [13]. There was no singular color-picking image source since RGB values for similar tissue varied due to differences in photographic parameters between images. Several CT datasets of healthy human subjects were volume rendered and displayed with this base color table. The interactive 3D volumes were viewed in stereo by several surgeons and the RGB values and interpolations were adjusted to obtain the most realistic images based on the surgeons' recollection from open, laparoscopic and thoracoscopic surgery. Specifically the thoracic cavity was examined with respect to the heart, the vertebral column, the ribs and intercostals muscles and vessels, and the anterior mediastinum. The abdominal cavity was examined with respect to the liver, gallbladder, spleen, stomach, pancreas, small intestine, colon, kidneys and adrenal glands. The result was a set of RGB values and a linear transparency ramp among them with the most realistic colorization and tissue discrimination via consensus among the surgeons. Table 1 displays the final values. Such iterative correction is to be expected as color realism is often a subjective perception based on experience and expectation.

Between the known tissue types, each RGB component value is linearly interpolated. Figure 4 graphically displays the realistic base color map values from Table 1 including the interpolated colors. In the case of the soft tissue, red primary color values are interpolated within the category. Simple assignment of discrete color values to each tissue type without linear interpolation produces images reminiscent of comic book illustrations. The lung appears pink, the liver appears dark red, fat tissue is yellow and bones are white, but there is nothing natural about this color-contrasted visualization. Inter-tissue interpolation produced the most

natural looking transition between tissue categories. More importantly, these linear gradients distinctly highlight tissue interfaces. Theoretically, as x-ray absorption is a non-linear function of tissue density, a corresponding non-linear color blending would be most representative of reality. However, there is no direct way to reverse calculate actual tissue density values from HU as important material coefficients in the CT radio density absorption model are not available. Future refinements via formal consensus development, multi-modality methods, or solving density absorption models, could assign specific colors to HU ranges at finer resolution or even to individual HU values.

5.2 Luminance Matching Conversion of Generic to Perceptual Color Maps

Conversion of a generically hued color map into a perceptual color map is accomplished by *luminance matching*. Generic color maps refer to those whose luminance does not increase monotonically as grayscale does, i.e., they aren't perceptual (see Figure 5). These are typical user selectable color tables (e.g. Figure 3 of our volume rendering controller shows choice of three generic color tables: Realistic, Spectral, Thermal). The thermal color table is sometimes referred to as a heated body or blackbody color scheme and is an approximation of the sequential colors exhibited by an increasingly heated perfect blackbody radiator.

Luminance matching takes advantage of the fact that grayscale is a perceptual color scheme due to its monotonically increasing luminance. Matching the luminance of a generic color map to that of grayscale for a given HU window will yield colorized voxels with the same perceived brightness as their grayscale counterparts. Eq. (2) describes how the luminance of hued color maps effectively becomes a function of HU, i.e., Y(HU). Luminance must be calculated using a color space that defines a white point, which precludes the HSV and linear, non-gamma-corrected RGB color spaces used in computer graphics and reported in this paper's data tables. We chose sRGB (IEC 61966-2.1), which is the color space standard for displaying colors over the Internet and on current generation computer display devices. Using the standard daylight D65 white point, sRGB luminance is calculated by Eq. (3):

$$Y(HU)_{color} = c_1 * R_{color} + c_2 * G_{color} + c_3 * B_{color} \qquad (3)$$

where: $c_1 = 0.212656$; $c_2 = 0.715158$; $c_3 = 0.0721856$ [14]

For a given HU window, grayscale values range from 0 to 255. Grayscale luminance is equal to the grayscale value (R, G, and B are all equal). If Y_{color} is greater than $Y_{grayscale\_}$, the linear RGB color values are converted to HSV. The value (V), or brightness component of HSV, is decreased, converted back to RGB, and luminance recalculated, repeatedly until the two luminance values are equal. If Y_{color} is less than $Y_{grayscale}$, there are two potential steps to increase Y_{color}. First V is increased, converted back to RGB, and luminance recalculated. If V reaches V_{max} (100%) and Y_{color} is still less than $Y_{grayscale}$, then saturation is decreased in similar stepwise fashion. Decreasing saturation is necessary because no fully bright, completely saturated hue can match the luminance of the whitest grays of the grayscale color map. Once the Y values match, the matched converted RGB values are used to build the color map for the specific HU window.

One of the advantages of using a realistic color map such as that shown in Table 1 is that colors always map to the same HU values from the full range of HU values regardless of the width of the imaged window (Figure 5). With the addition of luminance matching, the luminance values of all of window's colors seamlessly scale between 0% and 100%. This allows the greatest degree of luminance contrast for all ranges of window widths. For example, a small window centered on the liver will display reds and pinks with small, discernable, density differences made evident by perceptual luminance scaling. However, if a large window

centered on the liver is selected, the liver will appear dark red and will be distinctly identifiable relative to other tissues due to differences in both the display color and luminance. Grayscale, spectral, and thermal tables, on the other hand, dynamically map to the match the width of the selected HU window.

Figure 6 demonstrates the impact of realistic color and luminance matching. Each image was generated from the base grayscale visualization by selecting a generic color map and moving the Perceptual Contrast slider to the left or right (see Figure 3). Figures 6(b) and 6(c) should appear to be the most surgically realistic colorizations of the heart and surrounding anatomical structures. In contrast, the grayscale, spectral and thermal color maps, even in their perceptual versions, add a degree of distraction to the visualization due to their artificial colors. Most notable is that all the perceptual colorizations demonstrate the same appearance of depth and shape of the heart in the visualization, just as in grayscale (even without stereoscopy), whereas the non-perceptual ones vary (6(b) and 6(d) appear flat), even though some may add distinguishing features to the rendering. For example, in 6(b) the fat is artificially emphasized, and therefore more easily distinguished, relative to its appearance in the perceptual images (including grayscale) because of the high luminance of the yellow hue. Thus, for some uses, such as pulling out specific structures, thoughtful coloring that is perceptually incorrect can also be an advantage. From Figure 5 it should also be noted that the generic form of the thermal color map is already monotonically increasing in luminance, though not linearly like grayscale. It is thus understandable that 6(f) and 6(g) look nearly the same and that thermal maps are deemed almost as natural as grayscale by human users [11].

Figure 7 provides another example of the potential value of perceptually realistic color maps. Note how the hamstrings and gastrocnemius muscles surrounding the popliteal fossa are clearly delineated in Figure 7(a), but ill-defined in Figure 7(b). Luminance matching displayed in Figure 7(c) merges the perceptually desirable grayscale luminance of 7(a) with the clinically desirable realistic muscle colorization of 7(b) resulting in a visualization that exhibits the best of both color tables.

5.3 Mixing Perception and Reality

We have designed our GUI to allow us to *choose* the degree of realism and perceptual accuracy we desire for a particular color map via the Perceptual Contrast slider. This allows the user to view color maps in an arbitrary mixture of their generic or perceptual form. Alternatively, the user can choose to move the slider to only the end points, which represent generic color mapping (including anatomic realism) on the left and perceptual on the right. The slider mixes varying amounts of realism with perceptual accuracy by having Y_{color} match a linearly interpolated luminance as shown in Eq. (4).

$$Y(HU)_{interpolated} = (1.0 - P) * Y(HU)_{color} + P * Y(HU)_{grayscale} \tag{4}$$

$Y_{interpolated}$ is parameterized by the perceptual contrast variable $P \in [0.0 \text{ to } 1.0]$ and is the degree of mixing between generic and perceptual color mapping. The Perceptual Contrast slider on the GUI controls P's value. For any given P, $Y_{interpolated}$ is once again compared to $Y_{grayscale}$. The colorization algorithm once again dynamically calculates the HSV brightness and/or saturation changes necessary for the Y values to match.

The colorization algorithm further allows for sections of the anatomically realistic color map to overlap perceptual and generic color map values by selective exclusion of characteristic HU distribution regions. This is useful as realism is lost in some HU windows from luminance matching. For example, the fat color scheme tends to change from tan-lemon yellow to a murky dark brownish-green when it is in the lower range of a chosen HU window. Our interface has

checkboxes that exclude the fat region from luminance matching allowing it to retain its generic realistic color while the other regions display their color values with perceptual accuracy. Though this biases the visualization of the underlying HU voxel data, in some cases generic realistic fat colorization seems to improve the ease of interpretation of complex anatomy. In our implementation, the lung, tissue, and bone regions can also be selectively excluded from luminance matching.

5.4 Pseudocode for Perceptual Contrast slider using Luminance Matching

display User selected HU window with Generic Color Map

call Get_Perceptual_Contrast_Percentage_from_Slider

if Perceptual Contrast Percentage equals 100%

  print "Ymatch equals Ygrayscale. The generated color map will be perceptual"

else

  print "Ymatch is a linearly weighted mix of Ycolor and Ygrayscale."

for each HU Number in the User selected HU window

  call Get_HU_Number's_RGB_Triplet_from_Generic_Color_Map_1D_LUT

  call Calculate_HU_Number's_Ycolor_using_RGB_Triplet

  call Calculate_HU_Number's_Ymatch_using_RGB_Triplet

  call Convert_HU_Number's_RGB_Triplet_to_HSV_Triplet

  while Ycolor is greater than Ymatch

      call Decrease_HSV_Triplet's_Value_Component

      call Convert_HU_Number's_HSV_Triplet_to_RGB_Triplet

      call Calculate_HU_Number's_Ycolor_using_RGB_Triplet

  while Ymatch is greater than Ycolor

      if Value Component Value is less than 100%

          call Increase_HSV_Triplet's_Value_Component

      else

          call Decrease_HSV_Triplet's_Saturation_Component

      call Convert_HU_Number's_HSV_Triplet_to_RGB_Triplet

      call Calculate_HU_Number's_Ycolor_using_RGB_Triplet

  call Exclude_HU_Region (optional)

  save HU_Number's_RGB_Triplet_to_Luminance_Matched_Color_Map_1D_LUT

display User selected HU window with Luminance Matched Color Map

6. Discussion

Note that lightness (L) in the HLS color model is similar to value (V) in HSV space and should *not* be confused with L*. HLS lightness, much like HSV value and RGB brightness, is not a measure of the actual psychometric perceived brightness but rather a convenient approximation that does *not* account for the perceived difference in luminance among hues. Empirical

luminance matching of spectral hues with grayscale on non-photometrically calibrated monitors has been done with human users [15]. In these experiments, users are asked to match spectral hues with gray levels of sequential lightness in the HLS color space. Once again, facial images are used. Averaging the results from the users allows the authors to generate isoluminant spectral color maps and to manually generate hued color maps with increasing monotonicity. However, the method is neither automated nor real-time and requires a critical sample of human users to match the luminance for each uncalibrated monitor.

Generic anatomically realistic and spectral hued color maps (Figures 6(b), (d)) are not maximized for perceived brightness contrast and do not scale interval data with increasing monotonicity in luminance. For example, the perceived brightness of yellow in typical temperature maps is higher than the other spectral colors. This leads to a perceptual bias as the temperature data represented by yellow pixels appears inordinately brighter compared to the data represented by shorter wavelength spectral hues. A perceptual color map must mimic grayscale's monotonically increasing relationship between luminance and interval data value and maximize luminous contrast. Thus we have implemented these two perceptual criteria into our perceptual colorization algorithm.

Our methods theoretically offer greater insight into patient specific anatomic variation and pathology for teaching, pre-procedural planning, and perhaps in diagnostics and treatment. First, generically realistic colorization of three-dimensional volume rendered CT datasets appears to produce intuitive visualizations for surgeons affording advantages over grayscale. Second, our perceptual algorithm interactively generates color maps in real-time utilizing the standardized CIELAB color model whose parameters are empirically derived from human color vision research. Perceptually realistic colorization combines both the benefits of generically realistic colorization and the inherent perceptual attributes of grayscale producing color maps with the potential to enhance and correct the visual perception of complex, organic shapes.

Controlled task-specific user study and assessment, which are beyond the scope of this paper, would be required (e.g. in which diagnostic, planning, or treatment is performed via volume renderings) to verify the value and optimal balance of arbitrary discriminating colorizations and perceptually correct representation of data. Although fully perceptually corrected renderings could turn out to be considered impartial, a possibly key feature if one is concerned with ensuring diagnostic quality, it is likely advantageous to permit the user to also adjust the amount of perceptual correctness to "artificially" emphasize features, particularly under planning or treatment scenarios. We anticipate this to be the case in the fat region for the realistic colorization choices we've explored to date due to the high luminance of the yellow hue. Thus, we suggest use of a slider to enable each user to adjust the amount of "perceptual contrast" for any color visualization.

Note that CIELAB has competitor color models that claim to be more perceptually accurate. This does not devalue our luminance-matching algorithm since one could simply substitute a new luminance calculation from a more perceptually accurate color model to produce an improved perceptual color map. The colorization algorithm can also be extended to match non-monotonically increasing luminance distributions. For example, matching the desired luminance to an arbitrary, constant luminance value, i.e., $Y_{constant}$, easily creates isoluminant color maps. Note that in an isoluminant color scheme, $Y_{constant}$ is not a function of HU. We are not suggesting there are specific applications for isoluminant or other non-monotonically increasing color maps in medical imaging. Nevertheless, the colorization algorithm possesses this functionality should these color maps find their niche.

7. Conclusions

Our colorization algorithm draws upon feedback from surgeons analyzing reconstructed CT datasets on our 3D stereoscopic visualization environment combined with theory and empirical data regarding the recognition and generation of perceptual color maps by human users utilizing established results from the field of visual psychophysics. It appears that this merging of realistically colorized and perceptually accurate virtual anatomy insightfully interprets and impartially enhances volume rendered patient data for use in surgical planning.

As evidenced by the examples shown, grayscale color mapping of interval data can be automatically enhanced via the selection of realistic colors to represent the information in the available dynamic range and perceptually corrected via luminance matching to create intuitive visualizations. The generated color map gains realistic color while simultaneously maximizing luminance contrast discrimination in a manner typical of grayscale. Whether the color map spans the entire dataset or just a small subset (i.e. "window") of the dataset, the perceived brightness (L*) is spread from 0% to 100% luminance, thus maximizing perceptual contrast.

The colorization method has been applied to an interactive, 3D, volume rendered, stereoscopic, distributed visualization environment including real-time luminance matching of color-mapped data. The algorithm can be incorporated into other two- and three-dimensional visualization software or environments where color tables are used to represent information, as well as environments that render two- and three-dimensional representations of higher dimensional datasets, including in other domains.

Glossary

CIELAB (Commission Internationale de l' Éclairage 1976)

The CIELAB standard is a device-independent, perceptual color space that incorporates color-matching data from human vision experiments to exactly specify colors for a whole range of ambient lighting conditions. The details and geometric representation of CIELAB are outside the scope of this paper. CIELAB, or LAB for short, is comprised of L* (lightness), a* (spanning blue to yellow) and b* (spanning red to green) [16]. There are two points of interest regarding this color space compared to RGB/HSV:

1. RGB gamut, or extent of colors, is a subset of CIELAB's gamut. For example, CIELAB grayscale contains more "shades" of gray than RGB. As such, there are colors defined in CIELAB that CRT and LCD displays cannot produce [16].

2. CIELAB's geometric color space representation is calibrated to provide a perceptually uniform response by human observers that linearly correlate to the Cartesian distance separating two colors. For example, consider the L*, or lightness component axis. A color with twice the L* value will be perceived as twice as light. RGB and HSV do not possess this characteristic. In fact, outside of the line of grays, a color with twice the Value/Brightness of another in HSV space will most likely not be twice as bright.

D65

Illumination standard corresponding to a source emitting light at wavelengths identical to a perfect blackbody radiator heated to correlated color temperature of 6500K [14]. This is one of the recommended daytime illumination whitepoints in the CIELAB color space and is used by the luminance-matching algorithm in this paper.

HSV (Hue-Saturation-Value), HLS (Hue-Lightness-Value)

The HSV and HLS color models are non-linear transforms of the RGB color model [17]. Unlike RGB, whose mathematical simplicity is motivated by video hardware and computer architecture, HSV's components are designed to appeal to the intuitive manner in which colors are commonly described graphic artists (e.g. painters, printers, photographers etc). HSV is also referred to as HSB (Hue-Saturation-Brightness) where Value and Brightness are identical color space components. An inverted cone geometrically represents the HSV/HSB model where Hue is measured by its angle around the vertical axis (0 degrees for red, 120 degrees for green, 240 degrees for blue). Value is a measure of the amount of light, or hue intensity, and is defined by the height on the cone's vertical axis. V=0 (pure black) at the bottom point of the axis and V=1 (pure white) at the top point (the center point of the circle the is the cone's base). Thus the central vertical axis defines the same line of grays as the line diametrically spanning the RGB unit Cartesian cube from (0,0,0) to (1,1,1). Saturation is a measure of the amount of color, or hue purity, and its geometric measure is a radius line perpendicular to the vertical axis extending to the outer surface of the HSV cone. A hue is fully saturated (S=1) for a given value at the outer edge of the color space cone and desaturated (S=0) in the middle. For example, red de-saturates to pink as you head toward to center of the HSV cone when V=1. Saturation is also always zero at cone's bottom point because the radius is zero. Since HSV and HLS are based on RGB, they are also device-dependent.

RGB (Red-Green-Blue), RGBA (Red-Green-Blue-Alpha)

The RGB color model is motivated by the trichromatic theory of color perception that describes how human vision is most sensitive to RGB's three primary color components. RGB is commonly used in computer displays and utilized by graphics software applications where each component is described by an 8 bit binary representation [17]. RGBA adds a channel that is utilized by modern graphics processing units (e.g. computer video boards) to represent relative transparency. RGB is an additive color space (i.e. component colors may be linearly added to produce 2^{24} possibilities). RGB is a device-dependent since ambient illumination is not factored into the color model. As a result, a color represented by an identical RGB component triplet may look different on two un-calibrated displays or two calibrated displays in rooms with different ambient lighting environments. The RGB color space is geometrically represented by a unit cube in Cartesian coordinates with normalized component values of (1,0,0) for red, (0,1,0) for green, and (0,0,1) for blue and a line of grays diametrically bracketed by (0,0,0) for white and (1,1,1) for black.

sRGB (Standard RGB)

A non-linear transformation of RGB that is device-independent [14]. To calculate luminance matching, we must compute the CIELAB component L* which corresponds to a given RGB triplet. Calculation of sRGB luminance using RGB triplets is a necessary intermediate step in this process.

Acknowledgments

References

1. Pelizzari SA, Grzeszczuk R, Chen GT, Heimann R, Haraf DJ, Vijayakumar S, Ryan MJ. Volumetric visualization of anatomy for treatment planning. Int J Radiat Oncol Biol Phys 1996 Jan 1;34(1):205–211. [PubMed: 12118552]

2. Vidal FP, John NW, Guillemot RM. Interactive physically-based X-ray simulation: CPU or GPU? Stud Health Technol Inform 2007;125:479–481. [PubMed: 17377331]

3. Silverstein JC, Walsh C, Dech F, Olsen E, Papka ME, Parsad N, Stevens R. Immersive Virtual Anatomy Course Using a Cluster of Volume Visualization Machines and Passive Stereo. Stud Health Technol Inform 2007;125:439–444. [PubMed: 17377320]

4. Parsad N, Millis M, Vannier M, Dech F, Olson E, Papka M, Silverstein J. Real-time Automated Hepatic Explant Volumetrics on a Stereoscopic Three-Dimensional Parallelized Visualization Cluster [Abstract]. Int J CARS 2007;2:S488.

5. Bergman, LD.; Rogowitz, BE.; Treinish, LA. A Rule-based tool for assisting colormap selections. Proceedings of the IEEE Conference on Visualization '95; Oct 29–Nov 3, 1995; Washington D.C: IEEE Computer Society; 1995. p. 118-125.

6. Rogowitz BE, Treinish LA, Bryson S. How not to lie with visualization. Com Ph 1996 May/June;10 (3):268–273.

7. Mantz, T. Digital and Medical Image Processing [monograph on the Internet]. 2007 [cited 2007 August 28]. Unpublished Available from: http://www.cs.uu.nl/docs/vakken/imgp/ [>Literature >Chapter 2].

8. Stevens SS, Stevens JC. Brightness functions: Effects of adaptation. J. Opt. Soc. Am 1963;53:375–385. [PubMed: 13984028]

9. Poynton, C. Color FAQ [document on the Internet]. 2006 Nov 28 [cited 2007 August 28]. Available from: http://www.poynton.com/notes/colour_and_gamma/ColorFAQ.html.

10. Jackson, SA.; Thomas, RM. Cross-sectional imaging made easy. London: Churchill Livingstone; 2004. Introduction to CT Physics; p. 3-16.

11. Rogowitz, BE.; Kalvin, AD. The "Which Blair Project": A quick visual method for evaluating perceptual color maps. Proceedings of the Conference on Visualization '01; Oct 21–26, 2001; San Diego, California, USA. Washington D.C: IEEE Computer Society; 2001. p. 183-190.

12. Johnson, GM.; Fairchild, MD. Visual psychophysics and color appearance. In: Sharma, G., editor. Digital Color Imaging Handbook. Philadelphia: CRC Press; 2003. p. 115-172.

13. Ackerman MJ. The Visible Human Project. Proc. I.E.E.E 1998;86(3):504–511.

14. Lindbloom, B. 2003 Apr 20 [cited 2007 August 28]. Available from: http://www.brucelindbloom.com/ [>Math >Computing RGB-to-XYZ and XYZ-to-RGB matrices].

15. Kindlmann, G.; Reinhard, E.; Creem, S. Face-based luminance matching for perceptual colormap generation. Proceedings of the Conference on Visualization '02; Oct 27–Nov 1, 2002; Boston, Massachusetts, USA. Washington D.C: IEEE Computer Society; 2002. p. 299-306.

16. Ross, JC. The Image Processing Handbook. Vol. 5th ed. Florence: CRC Press; 2006.

17. Foley, JD.; van Dam, A.; Feiner, SK.; Hughes, JF., editors. Computer Graphics: Principles and Practice. Vol. 2nd ed. New York: Addison-Wesley Publishing Company; 1990.
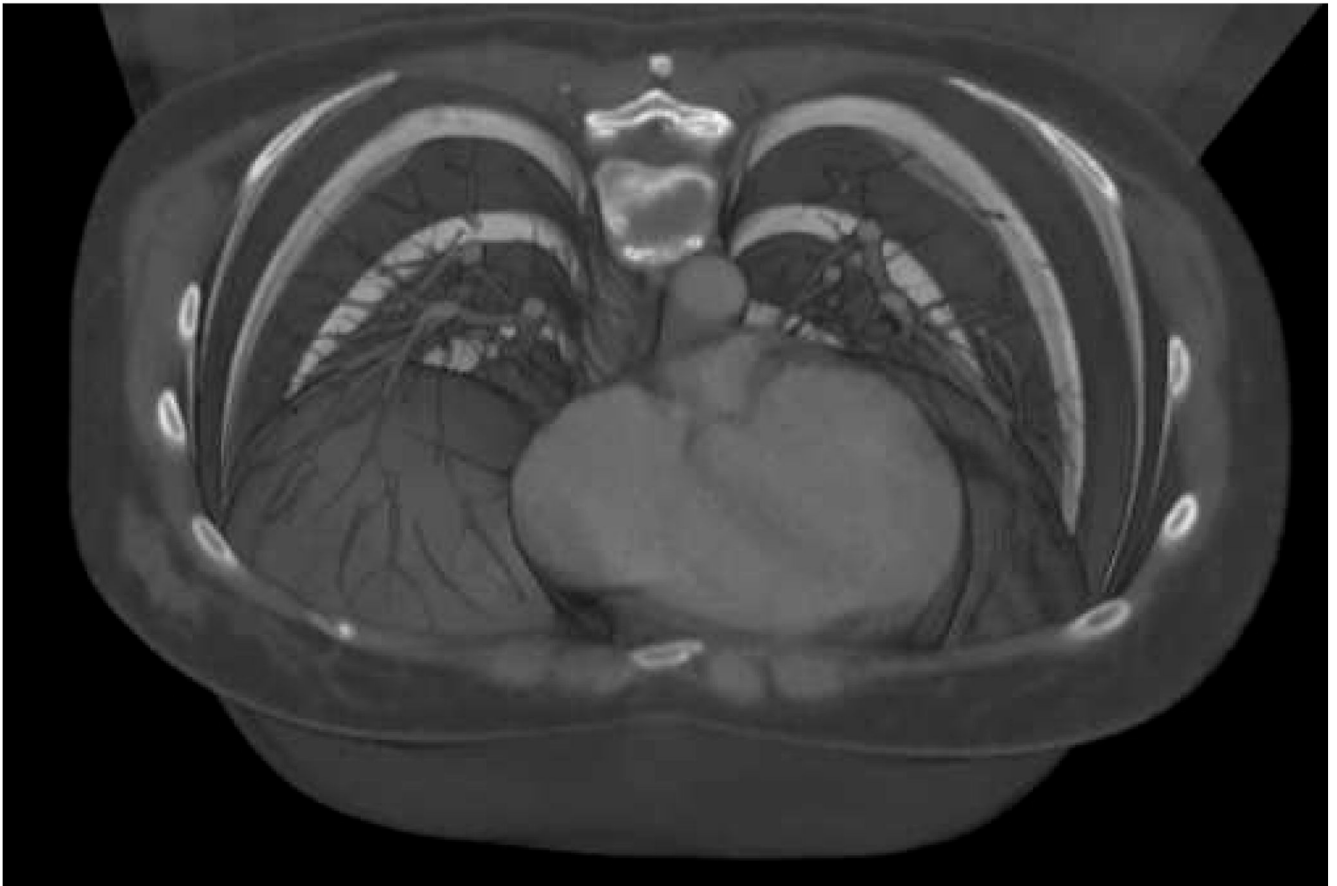
Figure 1.
Grayscale visualization viewed downward at the thoracic cavity in a bone window setting
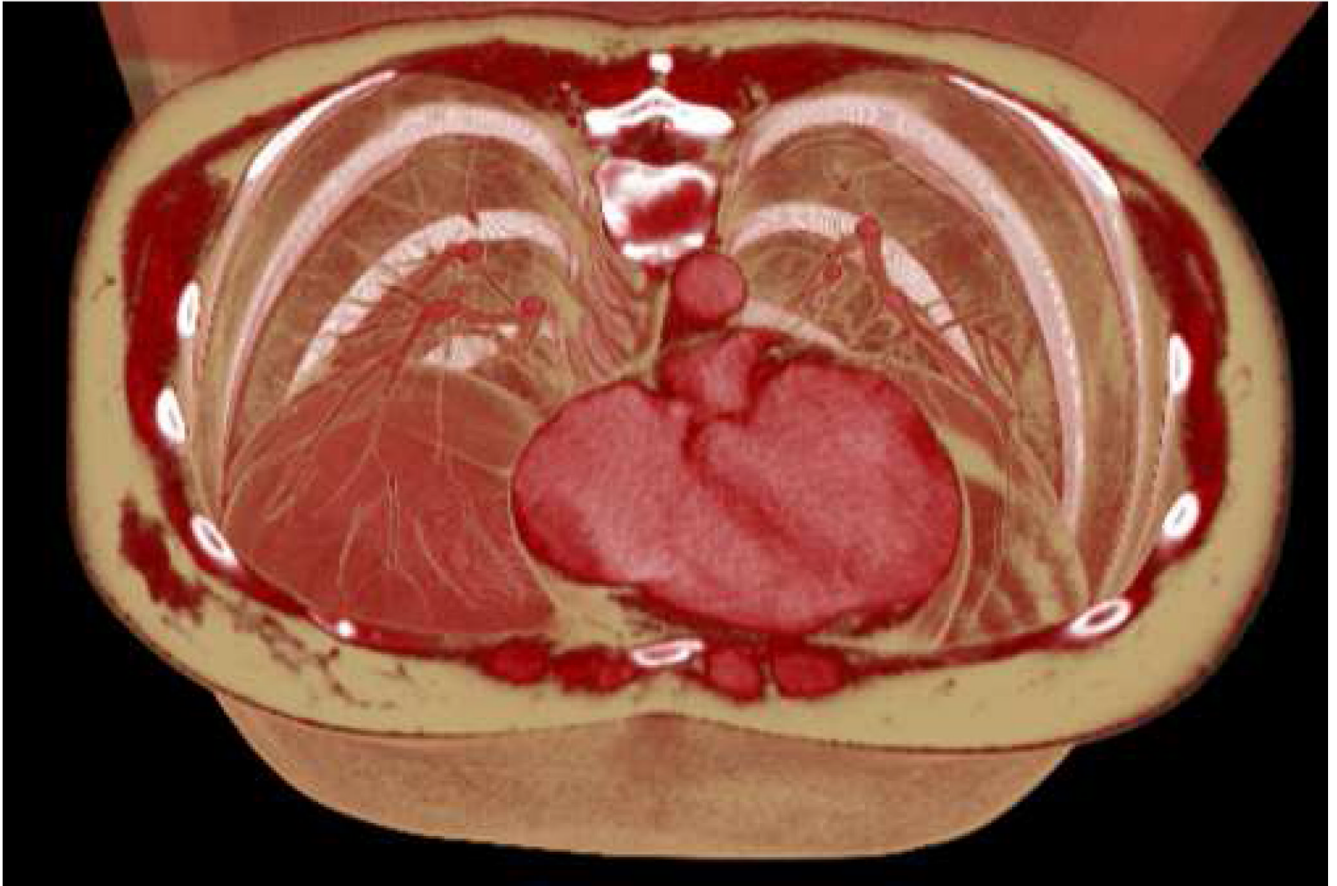(−400 HU to 1000 HU). The heart is seen as a dominant feature in the lower left-center.

Figure 2.
Identical visualization of the same dataset as in Figure 1, but using generic realistic color assignment defined by Table 1. The bronchia, diaphragm, heart, muscles and subcutaneous fat layer are more naturally delineated than in the grayscale colorization.

Figure 3.
User interface for volume visualization with color map parameters set as used in Figure 2.
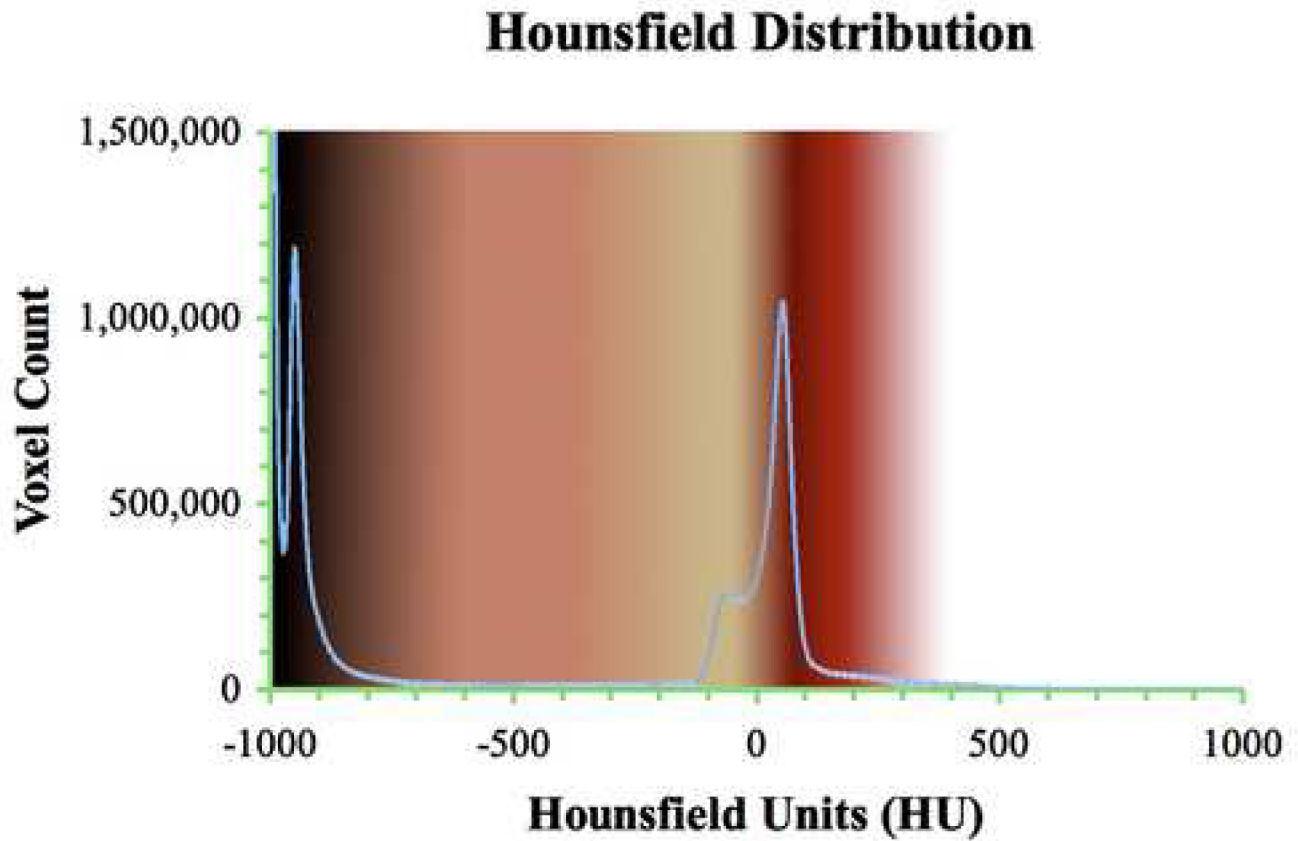
Figure 4.
Generic RGB values for realistic coloring including interpolated regions from Table 1. The method for assigning colors is described in section 5.1. The blue line is a superimposed HU distribution from a 1324 axial slice high-resolution CT dataset. The air, fat, and soft tissue peaks are apparent and the bone region appears as a long tail 600 HU wide.

Figure 5.
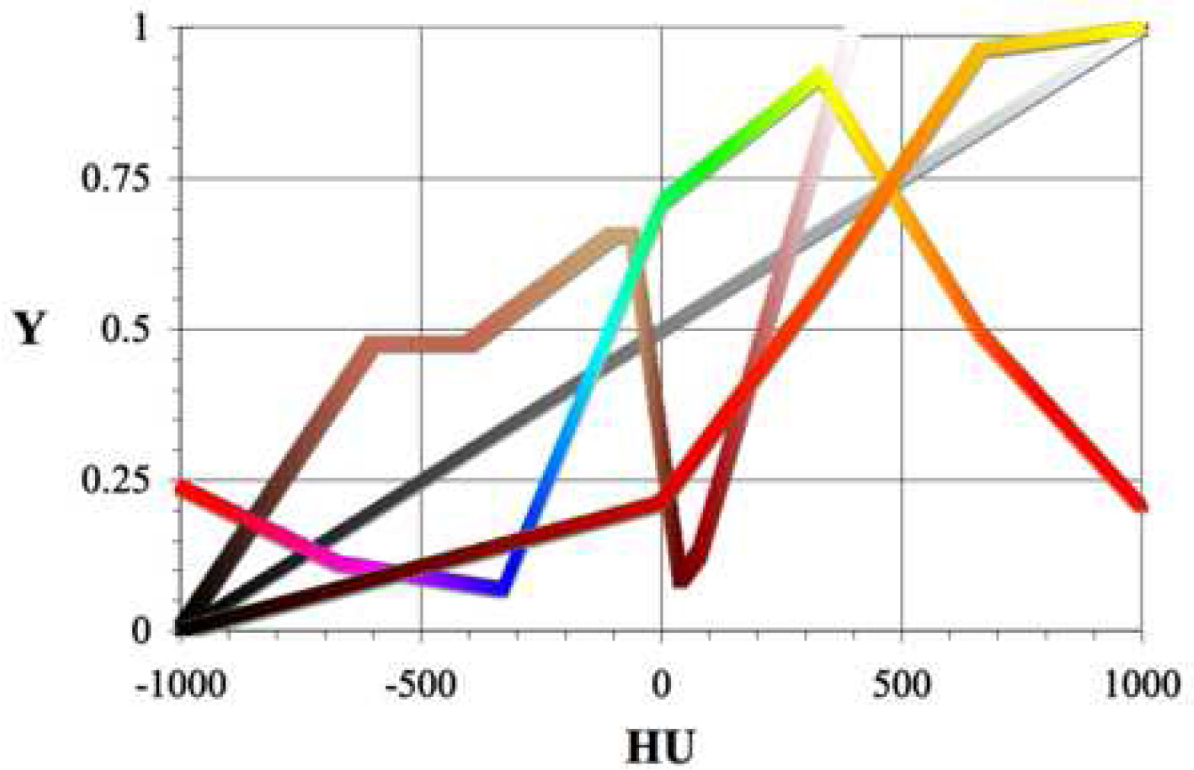Y(HU) for perceptual grayscale and the generic versions of the realistic, spectral, and thermal color tables over the full CT data range (maximum HU window). Typically, the grayscale, spectral, and thermal tables dynamically scale with variable HU window widths, i.e., the shape of the Y(HU) plot remains the same but the window's span of HU values determines the span of the plot. In contrast, realistic color schemes always map to the same exact HU values of the full HU window regardless of the selected window.
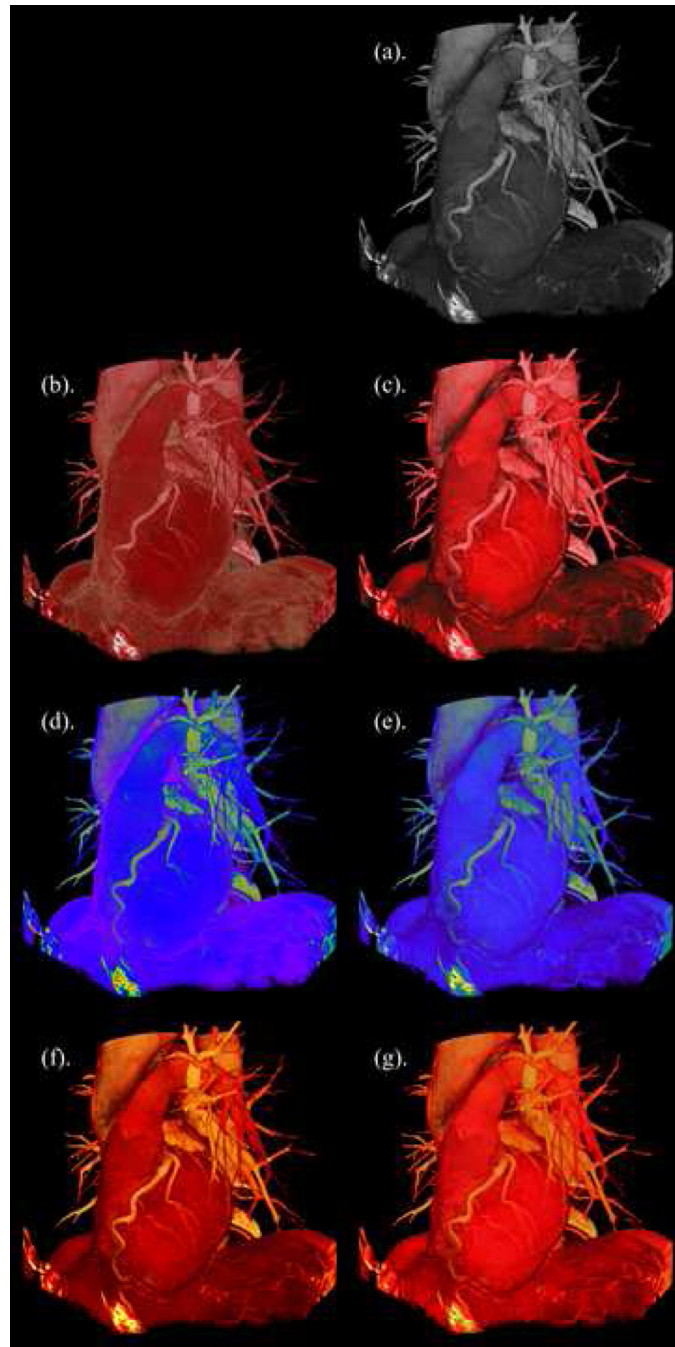
Figure 6.
Comparison of generic and perceptual color maps applied to CT volume visualization of the left side view of the human heart along with bronchi, vertebrae, liver and diaphragm. (b), (d), (f) are the generic versions of the realistic, spectral, and thermal maps respectively. (a), (c), (e), (g) are perceptual versions of the grayscale, realistic, spectral, and thermal color maps respectively.
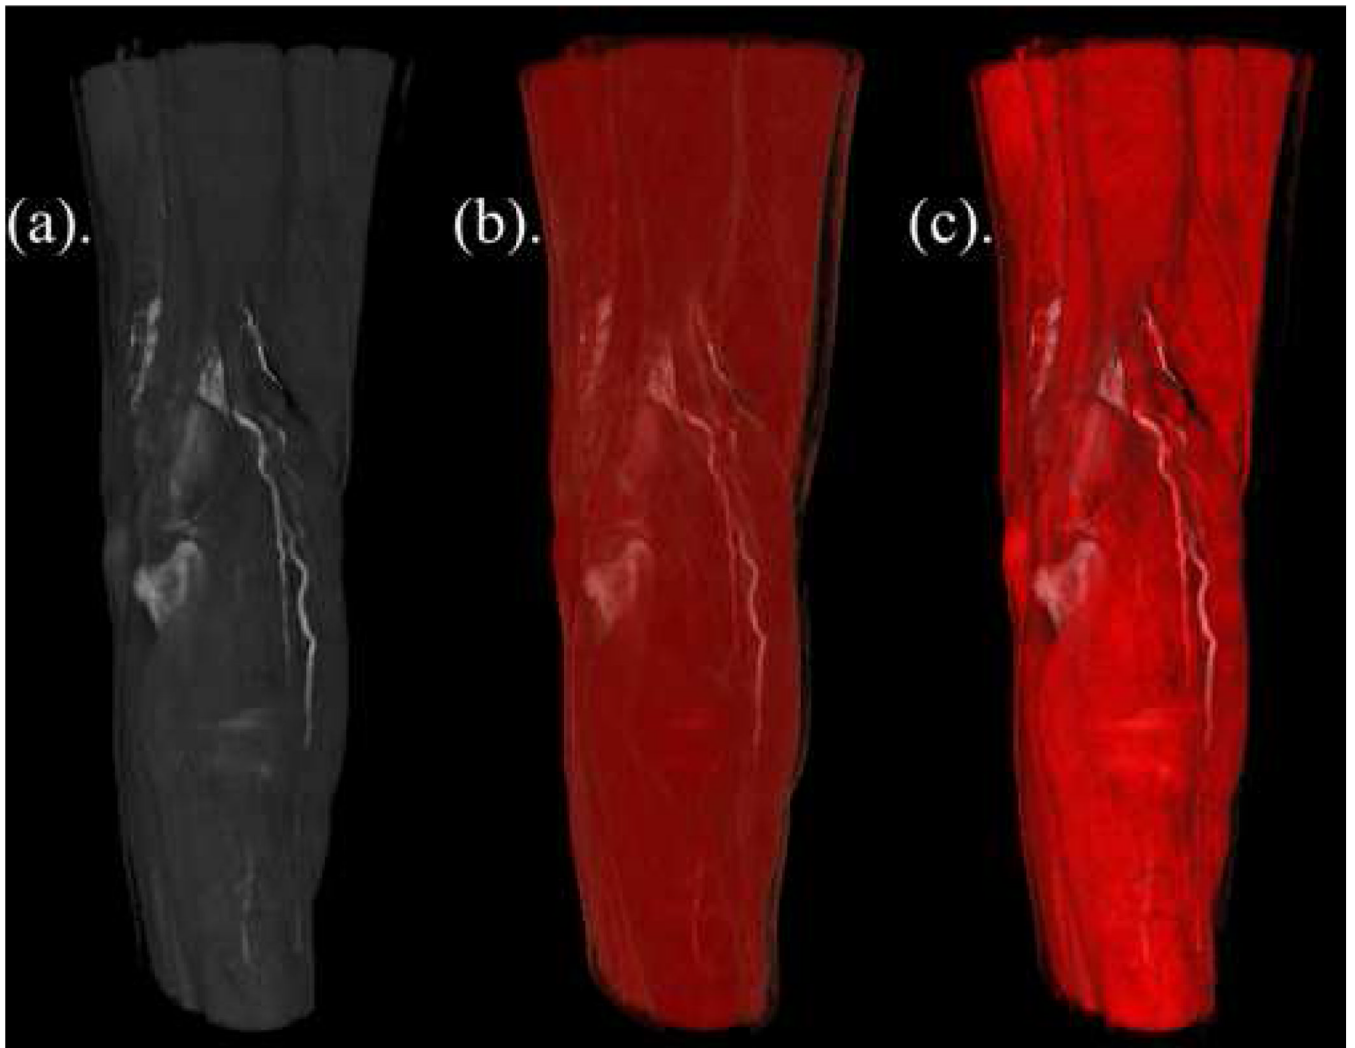
Figure 7.
Volume rendering from a CT scan with vascular contrast showing the posterior view of the knee including the hamstrings, popliteal fossa, and gastrocnemius muscles. From left to right, (a) grayscale (perceptual), (b) generic realistic coloring, and (c) perceptual realistic coloring.

Table 1

Generic Realism Color Table for Known HU Distribution Regions

Tissue type	HU Range	(R$_{base}$, G$_{base,}$ B$_{base}$)
Air filled cavity	−1000	(0, 0, 0)
Lung Parenchyma	−600 to −400	(194, 105, 82)
Fat	−100 to −60	(194, 166, 115)
Soft Tissue	+40 to +80	(102↔153, 0, 0)
Bone	+400 to +1000	(255, 255, 255)