The Colors of the Stars

Thor Olson
Management Graphics Inc.
Minneapolis, MN

Abstract

The brightness and colors of the stars have fascinated observers of the night sky for millennia. It is difficult however, to portray the depth and amount of visual information of the heavens on actual displays. This paper applies principles of reproduction, color science, and perception, to generate attractive and information-rich starfield images.

Although the casual appearance of stars is bluish-white, starlight actually ranges along the entire black-body locus. Scotopic (night) vision cannot see the full color that would result, but we can use color as a natural way to depict the various types of stars and their surface temperature. This is done in three steps. First, a relation between the astronomical color index of a star (the difference in power between two historically established spectral bands) and the star's effective surface temperature is obtained. From this black body temperature, colorimetric coordinates are calculated. Finally, the colorimetric values are converted to device coordinates for the display.

Introduction and motivation

Making pictures of the night sky is an ancient activity. Apart from astrological efforts to foretell the future, more successful uses for mapping the sky have been found for calendars, timekeeping, navigation, and deepening our understanding of the universe. And beyond these practical applications, star-filled skies find their way into the esthetic and spiritual domains of art and literature.

It would be useful for educational and reference purposes to make maps of the sky that portray the stars in an accurate way, their positions and relative brightness correctly displayed. Making such a technically accurate image is not possible with common display technologies however, and even if it were, its usefulness would be limited to about the equivalent of walking outside at night and looking up. The dynamic range of naked-eye stars is about 100:1, not a coincidental match to the human visual range for a given adaptation. This is a severe restriction; of the millions of stars we might want to show, only a few hundred could be accommodated.

It would be nice if star images could also convey something about the nature of their light. Although most stars appear to us as bluish white, the stars vary enormously in temperature, and consequently, color. Presenting stars in their correct color would provide a natural indicator of a correlated physical characteristic.

Although an exact duplicate of the night sky may not be practical or even useful, this is actually a rather familiar problem in the area of graphic arts and imaging. It is almost always the case that we do not want to duplicate a scene, but instead want to duplicate the *appearance* of that scene. We apply principles of reproduction, color science and psychophysics to accomplish this. This paper describes the application of those principles to make information-rich and natural-appearing, even beautiful, representations of the night sky.

Historical notes on star color and magnitude measurements

Before the invention of photography, stars were charted according to their relative visual brightness. Astronomers became skilled at making judgments of brightness by comparisons with reference stars. A scale of measurements evolved where the brightest stars, those "of the first magnitude" were judged to be twice as bright as second magnitude stars, which were twice as bright as third magnitude stars and so on. The limit of naked-eye astronomy was around magnitude 6. Much later, when instruments could measure actual stellar luminances, it was discovered that a difference of 5 magnitudes was not 32:1 as expected, but was instead 100:1. This means that a stellar magnitude unit is equivalent to the 5th root of 100, or a factor of about 2.5. (This historic anecdote provides a clue regarding the visual brightness response of scotopic vision).

When photographic films became sensitive enough, they were put to the task of automating the measurements of star magnitudes. It became immediately apparent however that the magnitudes as recorded on film did not match the magnitudes that had been so tediously obtained by observers over the years. The discrepancy was traced down to the difference in spectral sensitivity between the film (which in those days was blue sensitive) and the scotopic sensitivity of humans.

Even though there was not a perfect correlation, making magnitude measurements with film was so vastly more efficient that stars were catalogued according to their "B" (blue) magnitudes in addition to their "V" (visual) magnitude. It was clear that B and V were samples of a continuous spectrum from each star, and the difference B-V became known as the star's color index.

Converting color index to black body temperatures

The spectra of the sun and stars is very complex, with emission and absorption lines that indicates the chemistry and physics going on at their surfaces and in their atmospheres. But underlying the complexity is a continuous spectrum that can be approximated as an idealized black body radiating at the star's characteristic temperature. The temperature of stars ranges from 3^0K (if you include cold dead stars) to over $50{,}000^0$K. The spectral power radiated by them follows laws discovered by Max Planck at the beginning of the century. A sampling of the relative spectral power output for black bodies at various temperatures is shown in figure 1.

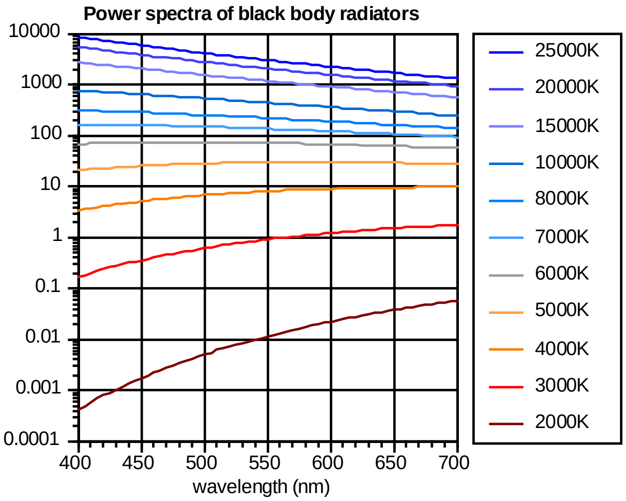

Figure 1 Spectral power output for black body radiators at various temperatures.

The spectral differences that led to B and V magnitude measurements was formalized by astronomers. Using technology available at the time, specialized glass filters were selected that approximated the spectral sensitivities of photographic film and human visual response. Plots of them are shown in figure 2. If you were to make measurements of stars through these filters you would find that their ratio (or *difference* in magnitude units), provides an indirect measure of their temperature. The use of "magnitude" as a unit was preserved, and even though these are actually radiometric measurements, astronomers refer to this activity as stellar photometry. The difference between the B and V magnitude measurements became known as the stars color index.

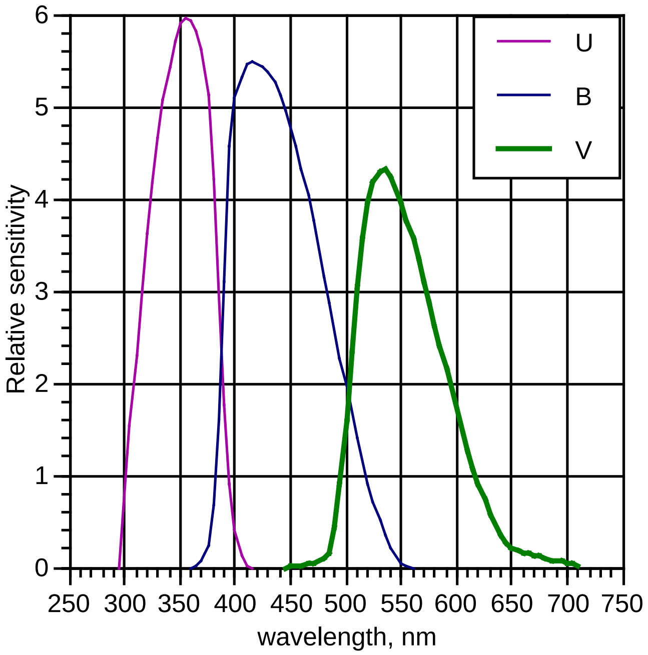

Figure 2. Spectral characteristics of the U, B and V bands used in astronomy.

An exact formula relating the color index, B-V, to the effective black body temperature is not known to this author, but the energies transmitted through the B and V filters can be computed, and effective temperatures plotted for various luminosity classes of star spectra (figure 3). An approximating formula was obtained to fit this data. From it one can find the effective temperature of a star given its color index.

$$T_{eff} \approx 1000 + \frac{5000}{(B - V + 0.5)} \qquad (1)$$

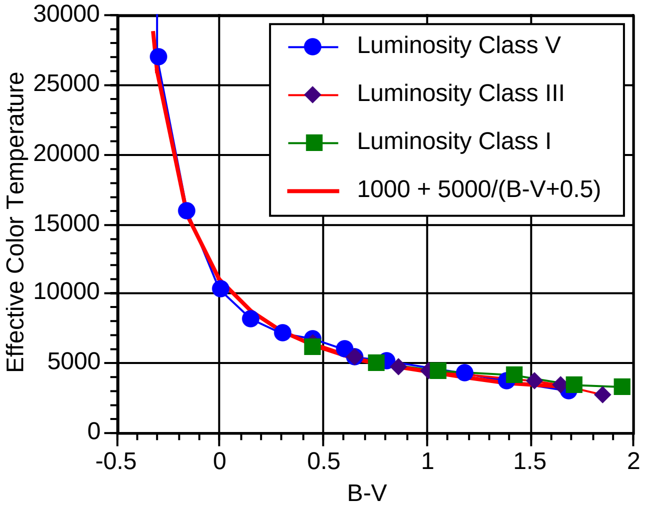

Figure 3. Effective black body temperature as a function of color magnitude index B-V

Converting temperature to color

Just as we can compute the energy through the B and V filters, we can obtain the energies sensed by the CIE colorimetric standard observer and obtain chromaticity coordinates for various black body temperatures. In fact, this was done long ago and the locus of black body radiators is shown on many chromaticity charts as a reference (figure 4). A polynomial fit to the locus is provided in Hunt [1] which can be used to obtain the chromaticity of a star from its temperature.

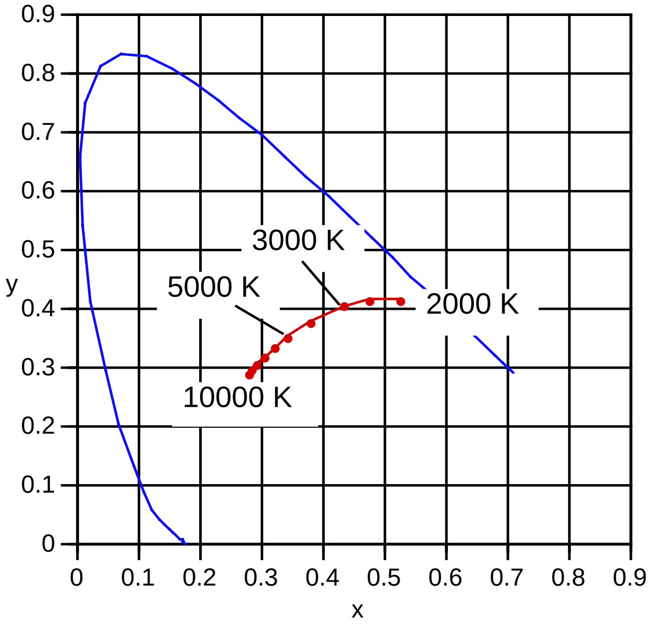

Figure 4. Locus of black body radiators on an xy chromaticity diagram.

Converting color to device coordinates

We are now faced with how to present the color on a real output device. Up to this point the conversion from stellar magnitude measurements has followed well defined operations. We will now have to make some judgments, guided by principles of perception, on how best to present this information on a physical device. Two types of output will be discussed: a CRT monitor, and film transparencies.

CRT displays

A CRT is a well behaved color device. Once it has been calibrated and characterized, the colors it makes are predictable and easy to obtain. Say that a monitor has been set up so that its white level matches the D65 daylight whitepoint. If it is using the standardized phosphor set in the CCIR Rec. 709 video standard, then the linear RGB levels can be obtained from an XYZ triple by the following matrix operation:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.240 & -1.537 & -0.499 \\ -0.969 & 1.876 & 0.042 \\ 0.056 & -0.204 & 1.057 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \qquad (2)$$

The actual CRT drive levels R'G'B', are obtained from RGB by applying the power law of the device which, when calibrated to the video standard, uses a gamma of 0.45 (approximately 1/2.2).

$$\begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} = \begin{bmatrix} R^{0.45} \\ G^{0.45} \\ B^{0.45} \end{bmatrix} \qquad (3)$$

One may now take the chromaticities along the black body locus representing stars at those temperatures, scale them to obtain XYZ tristimulus numbers with Y=1 (in other words, treat it as if it was an adaptation whitepoint), and calculate video R'G'B'. This has been done for a collection of temperature points along the locus and is plotted in figure 5. This plot tells us what monitor R'G'B' levels to use to display stars at a given temperature.

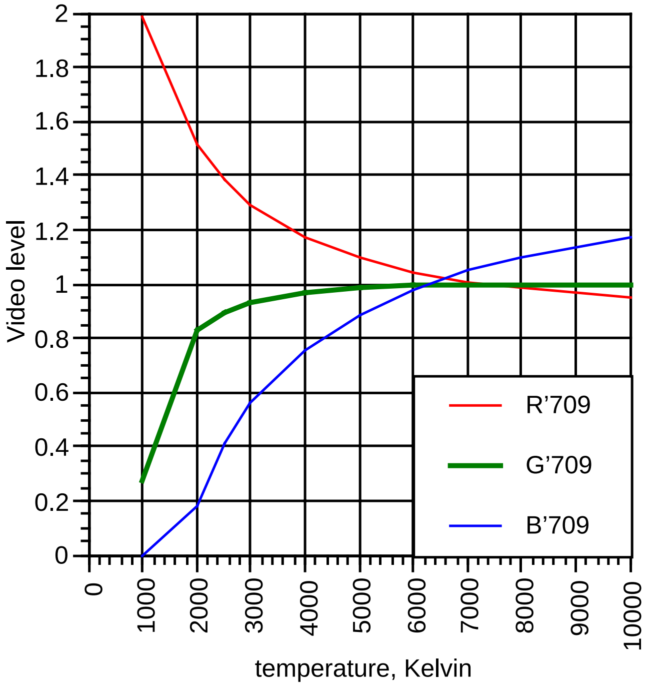

Figure 5. Video R'G'B' to represent star color as a function of temperature.

There are some interesting things to point out in this plot. First, as astronomers know, there are no green stars! At no temperature does one find that the energy is highest in the green band. Either red is dominant or blue is dominant. Even at the point where they are equal, they are both greater than the contribution from green.

The green channel stays relatively constant. By the selection of our device whitepoint (D65) green equals unity at or near 6500^{0}K. At that point red and blue are both greater than one. What this means is that the black body locus does not pass exactly through the D65 whitepoint (if it did, all three components would equal 1 at 6500^{0}K).

There is another practical matter to consider. Figure 5 shows the monitor levels normalized to 1, usually taken to be the maximum drive level in each color channel. In order to accommodate the full temperature range, we must scale the levels so that no color channel will run out of range. The scaling must be done in RGB prior to its nonlinear conversion to monitor R'G'B'. A factor of 2 would be enough to display stars down to nearly 3000^{0}K or so. But this would cause the near-white stars at 6500^{0}K to be displayed at a rather dim level, considerably down from white.

There are several approaches to solving this. One is to simply accept the dim white point, knowing that the relative color strengths are correct. We could also scale the levels so that the maximum of each RGB set was normalized. This will keep all of the temperatures within range, and they will be chromaticity-correct, but the relative perceived brightness of differently colored stars will be skewed. Low temperature red stars especially will seem dimmer than they should (remember, our sample star colors were scaled so that their luminance level, Y, started out the same). From an esthetic viewpoint, suffering this brightness error is preferable to an overall dim picture. Fortunately there is a way to restore the brightness, described later in the treatment of actual star magnitudes.

Film images

Compared to CRTs, film is considerably less well-behaved. Although it has a large dynamic range and color gamut, there are variations in film batches and processing chemistry. On top of this are differences in projector bulbs and room lighting. Understanding the colorimetry of projected transparencies has been tackled by several researchers who continue to discover its complexities [2][3][4]. Nevertheless, because of the drama and depth it adds, and the natural similarity of a darkened theater to a nighttime environment, it is a prime output medium for this project. If we are willing to sweep the uncontrolled variations under the observer's adaptation rug, we can still obtain perceptually satisfying images of the stars.

The range of star colors is not so high as to strain the limits of color slide film. In fact, they are relatively unsaturated colors, and because of this, the linear matrix type operation is successful. The matrix for Ektachrome film exposed on a Management Graphics Solitaire film recorder is:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.275 & -1.557 & -0.511 \\ -0.781 & 1.693 & 0.045 \\ -0.030 & -0.119 & 1.054 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \qquad (4)$$

Film recorders are typically calibrated to a power law with a gamma of 2.2, so there is the same nonlinear transform to R'G'B' as for CRT monitors.

Measurements were made on some test "stars" made from the formulas above and imaged onto Ektachrome film. The samples are illustrated in figure 12 and their chromaticity measurements plotted in figure 6. The success in the match with the targets indicates that film can be treated as a linear system over the small range that these black bodies cover.

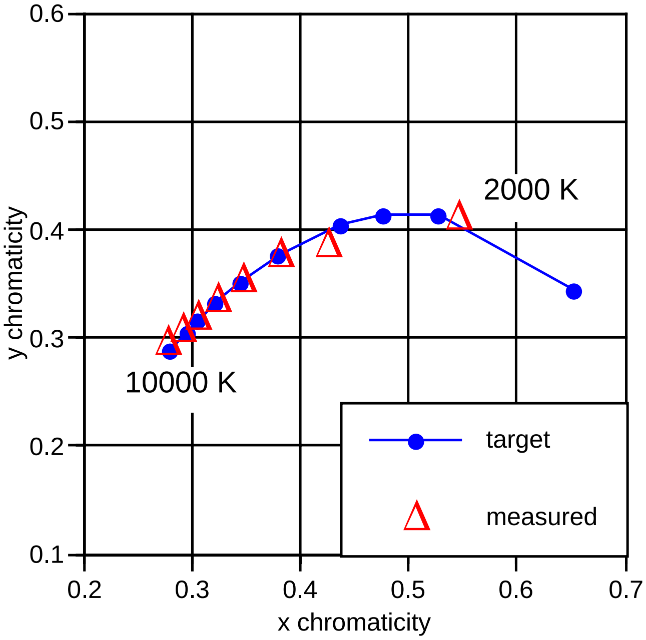

Figure 6. Measured colors from Ektachrome

Perceptual mapping of stellar magnitudes

So far we have discussed only the color of stars. The other major characteristic of the night sky we would like to portray is the enormous range in their apparent brightness. The naked eye can see stars down to magnitude 6 (on a clear dark night). But the sky is filled with stars dimmer than this that show up when you look closely with binoculars and telescopes. Can these be shown in a meaningful way in a star field image? Further, and more basic, stars are essentially point sources. How does one represent a point source on a output device that uses finite-sized pixels?

As mentioned in opening, maps of the sky have been made for a long time. Modern star maps have settled on various conventions to display the brightness and types of stars. A look at one of the maps of sky cartographer Wil Tirion [5] shows the common practice of using larger diameters to represent brighter stars. For a black and white chart, this is really the only option available. It has a basis however, in how stars are recorded photographically, and are perceived visually.

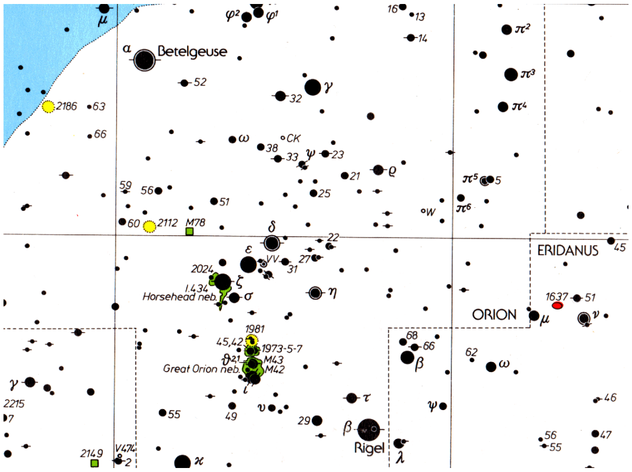

Figure 7. A section from a Tirion star map.

Before electronic photometers became available, and even now when making large surveys, star magnitudes were obtained by measuring the diameter of their recorded image on film. The principle behind this is shown in figure 8. The star's image is not a point, but is the point spread function of the lens viewing it. This function will depend on the diffraction limit of the aperture, and the aberrations of the optics, but usually has most of the energy localized at the center and then trailing "skirts" farther away. The behavior of film is to convert to metallic silver those areas that are exposed with enough energy to exceed a certain threshold level. Because of this, brighter stars, having a higher amplitude point spread function, will make a larger recorded spot. The relationship is not proportional, but depends on the details of the point spread function of the system, and the threshold of the film.

The human eye does not behave like film, but does share some characteristics. The diffraction and aberrations of the lens produces an optical point spread function, and the receptors in the retina have thresholds that cause neuron activity. Unfortunately, there is a great degree of retinal processing that prevents a simple relation between luminance and perceived brightness. Much research has been conducted into this topic, but it is largely limited to experiments using *area patches* of known luminance, not point sources and their equivalent area stimuli.

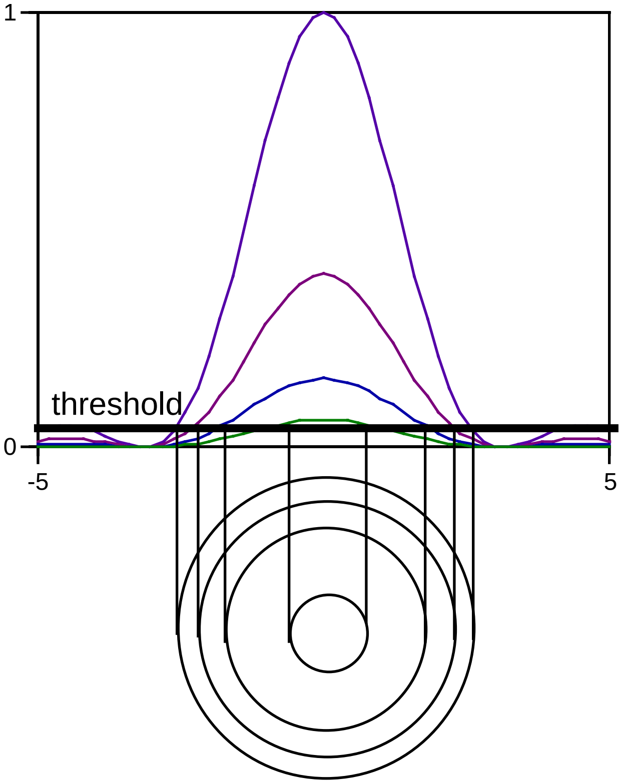

Figure 8. Recorded spot sizes on film for various intensities (a four magnitude range) in a diffraction limited system.

The equal energy hypothesis

In lieu of the missing psychophysical data, some experimentation was done. We would like to obtain a scale of small *areas* that produce that same relative sensations as their corresponding point sources. An early hypothesis translated the intensity of a star directly into an area on the image. This would preserve relative energy, a magnitude 5 star would have 2.5 times the area of a magnitude 6 star. But making the area of a star image proportional to its luminance was disastrous. The range in star sizes became enormous, obviously not corresponding to a perception of the night sky (figure 9).

The equal sensation hypothesis

Making the area proportional to perceived *brightness* on the other hand was quite believable. Brightness is asymptotic to the 3rd root of luminance, so the growth in area was considerably smaller, $\sqrt[3]{2.5}$ or 1.36 per magnitude unit. This strategy makes some sense. Consider a point source that is 1 magnitude (2.5x) more luminous than another. If the intensity of the sources are beyond the adaptation of the eye, then according to the CIE functions [6], it will be perceived as 1.36 times brighter. If one wants to produce the same total *sensation* as this, but with a fixed luminance within the eye's current adaptation, an area that is 1.36 times larger seems reasonable.

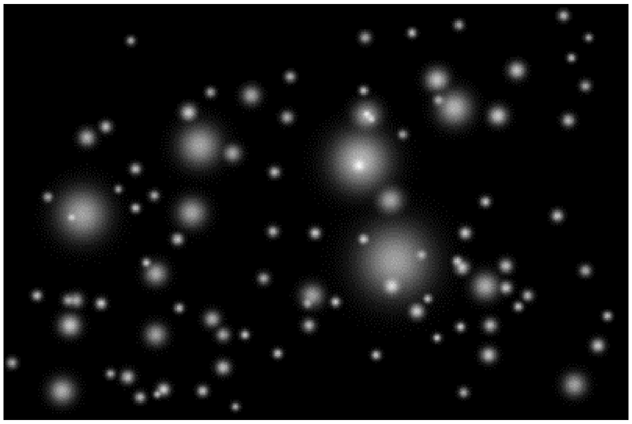

Figure 9. An unrecognizable constellation Leo, rendered using the equal energy hypothesis.

This worked well for bright stars which had visibly distinguishable areas, but it caused difficulties in discriminating the dim stars that were imaged using very small areas. It seems that once the retinal image area approaches the eye's point spread function (where small areas look the same as point sources) one needs to return to the principle of delivering similar energies in order to achieve similar brightness perceptions.

A single formula that behaves in the desired way at the extremes, and makes a nice transition is:

$$R^2 \propto \frac{L}{\left[L^{2/3} + L_A\right]} \qquad (5)$$

where R is the radius of the displayed star, and L is the relative luminance of the star obtained from its visual magnitude M:

$$L = 100^{\frac{-M}{5}} \qquad (6)$$

L_A is a transition luminance, corresponding to a transition magnitude where the area of the displayed star changes from being proportional to brightness ($L^{1/3}$) for bright stars, to being proportional to luminance, for dim stars. The selection of this transition point will depend on the visual angles covered by the star images.

This produces the relation plotted in figure 10 where the transition magnitude was M=2. This plot also shows the scale used in Tirion's star map. It is not known how he devised the scale, but it is satisfying that it is similar to our heuristic derivation.

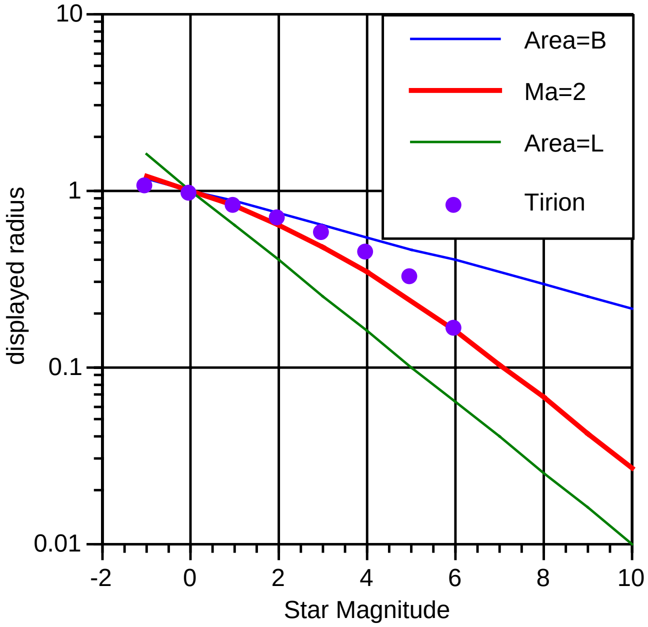

Figure 10. Relations used to draw stars using area as an indication of their magnitude.

The limits of area equivalence

In continuous tone images, we have the luxury of being able to reduce the intensity of the star in addition to controlling its rendered size. Over a large range of star magnitudes, the relative area between the brightest and dimmest star becomes quite large. If the dimmest star is rendered using a single pixel, it is possible that the brightest star could become excessively big, consuming too large a fraction of the display area. To solve this, the brightest star is set to an "esthetically acceptable" diameter (this judgment will depend on the purpose and resolution of the image). All dimmer stars are then smaller than this. Eventually a star magnitude will be encountered which would require less than 1 pixel. We can now reduce the amplitude of the pixel to preserve the brightness relation. Instead of the amplitude being fixed and the size reduced, we now convert to keeping the size fixed (at more than 1 pixel if we desire) and reducing the amplitude. We can continue this until we have reach the dimmest level of the display device, at which point we have reached the limit of stellar magnitudes that we can display in their proper relative brightness.

The reduction in amplitude for each magnitude should be the same as the reduction in area. Since this presumably occurs for small dim stars where the area was proportional to luminance, a factor of 1/2.5 = 0.4 is needed. Since the image display follows a power law, the reduction in drive level for each magnitude step is only 0.66. The combination of using area and amplitude was used to make the magnitude scale shown in figure 11 (which may not reproduce to its full range in this publication).

Recovering the displayed color range

The star magnitude to size and amplitude relations can now be put to use to solve the problem of displaying the full color range of the stars. The difficulty was that to show the correct colorimetric range, we had to accept either a brightness error, or a dim white point. We can now restore the whitepoint to its normal full level, normalize star amplitudes for their color, and adjust the displayed size of the star to preserve relative brightness. For example, a low temperature red star would be amplitude normalized to its red channel, reducing the green and blue. This reduces its visual magnitude (brightness). The perceived brightness is restored by increasing the area used to display that star.

In effect, we have made a tradeoff. We are obtaining an improvement in the perceptual presentation of the stars, but the calibration between magnitude and displayed area now depends also on the color. The resulting color scale is shown in figure 12, where the adjustments in size can be observed to preserve an overall impression of similar brightness.

Additional elements in starfield illustrations

Almost by definition, a good map is an effective presentation of information. A poor map may contain the same data but presents it in a way that makes it difficult to understand. The choice of scales, symbols, colors, and fonts are critical to the success of the map. There are now a number of "desktop planetarium" programs, and most will make star charts of a selected area of the sky. In some, the user is free to make color choices for symbols, coordinate grids, and backgrounds. It is observed by this author however, that any choices much different than the traditional colors tend to make it *more* difficult to read the chart.

There is probably a large body of cartographic knowledge that helps a professional mapmaker to make these color choices. It is interesting that the most effective map is also an attractive map. There are likely esthetic and psychophysical reasons for this, but it is beyond the scope of this paper to explore that direction.

Additional elements that make a rich map of the sky include gridlines showing the celestial coordinates, an indication of the faint band across the sky made by the Milky Way, and sometimes, construction lines showing the constellations.

A natural background for an image of the night sky is black of course. Although difficult to print, it makes the stars stand out in excellent contrast in a transparency. The selection of gridline colors is made so that the grid is an unobtrusive element, almost invisible, until one wants to see it, at which point is should be visible for unambiguous position reference.

The Milky Way is an awe-inspiring feature of the night sky. The faint glow of the band of our galaxy is not really *seen*, but rather *sensed* by the low light sensitivity of off-axis rod cells. One has the feeling that it is a bluish white glow. This is due to the high percentage of young hot stars (blue stars) in the nearby arms of our galaxy. Rendering the Milky Way at the infinite temperature limit for black body radiators provides one more color reference, albeit somewhat artificial, for making a visually appealing map.

Constellation lines are a dilemma. Like the coordinate grid, one wants them to be available to help locate these major star patterns. But once found, it would be nice if they could disappear completely, leaving the beauty of the starfield to make its own patterns in the mind of the viewer. If one includes the constellation lines in an illustration, they can be rendered in a faint, transparent color, say greenish, that contrasts with the star colors and the Milky Way.

Combining all of these elements into the color and brightness-scaled star fields yields the typical view shown in figure 13, a satisfying and informational image.

Beyond the flat map

There is another dimension that can be utilized to present information from an image: depth. A stereo view of the sky would normally be perceived as the sky looks: flat. The stars are just too far away to provide any significant parallax. But there is no reason we cannot amplify the small parallaxes that have been measured by astronomers to illustrate the relative *distances* of the stars we see. This has been done in figure 14. If one can fuse the two images in the pair, either through free-viewing or using a stereo viewer, one can get a glimpse of how God might see this part of the sky (provided God has trichromatic night vision and an interocular distance of 1.6 light years)!

More questions than answers

As usual in investigations of Nature, as questions are answered, more are posed. If we can display the colors of stars, and a 6500^{0}K star looks white, why does the Sun (a 6800^{0}K star) look yellow? And why is the sky blue, really? Shouldn't the sky scatter starlight as it does sunlight? Why don't stars look red when they set? The images created here are not calibrated views of the sky. Nevertheless, for educational and artistic purposes, it seems that it is worthwhile to pay attention to the principles of esthetics, color, and perceptual sciences.

References

1. R.W.G. Hunt, Measuring Colour, Ellis Horwood Limited, West Sussex England, 1987.
2. Robert Hall Wallis, *Film Recording of Digital Color Images*, USCIPI Report 570, May 1975.
3. Mark Fairchild, Roy Berns, Audrey Lester, and Hae Kyung Shin, *Accurate Color Reproduction of CRT-Displayed Images as Projected 35mm Slides*, IS&T/SID Color Imaging Conference, Scottsdale AZ, 1994.
4. Audrey Lester and Mark Fairchild, Thermochromism of Ektachrome 100 Plus Professional Transparencies Upon Projection, Journal of Imaging Science and Technology, v.38, n.4, July/Aug 1994.
5. Wil Tirion, Cambridge Star Atlas 2000.0, Cambridge University Press, Great Britain, 1991.
6. H.W. Bodmann, P. Haubner, A.M. Marsden, *A Unified Relationship between Brightness and Luminance*, CIE Proceedings 19th Session, Kyoto Japan, 1979.
7. Elske v.P. Smith and Kenneth C Jacobs, Introductory Astronomy and Astrophysics, W.B. Saunders Company, Philadelphia PA, 1973.
8. Tom N Cornsweet, Visual Perception, Harcourt Brace Jovanovich, Inc.Orlando FL, 1970.
9. David H Levy, Skywatching, The Nature Company, Berkely CA, 1994.
10. David Malin and Paul Murdin, Colours of the Stars, Cambridge University Press, London, 1984.
11. M.J.G. Minnaert, Light and Color in the Outdoors, Spreinger-Verlag, New York, 1974.

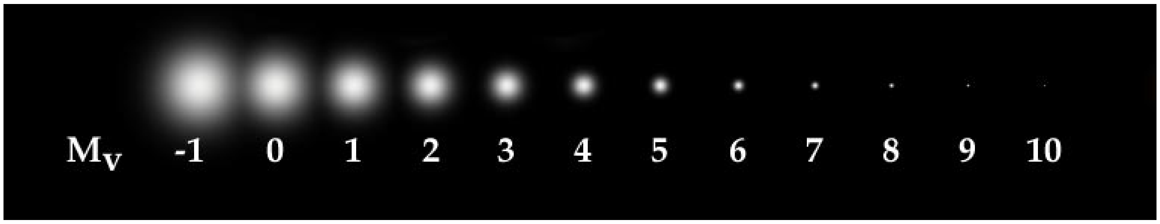

Figure 11. Rendered magnitude scale using both area and intensity to extend its range.

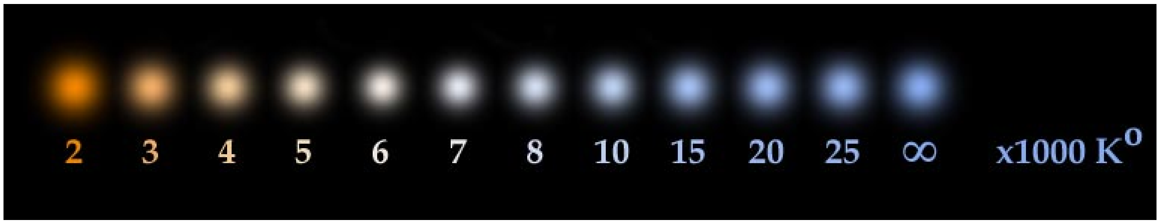

Figure 12. Rendered color scale using area variations to preserve brightness at different temperatures.

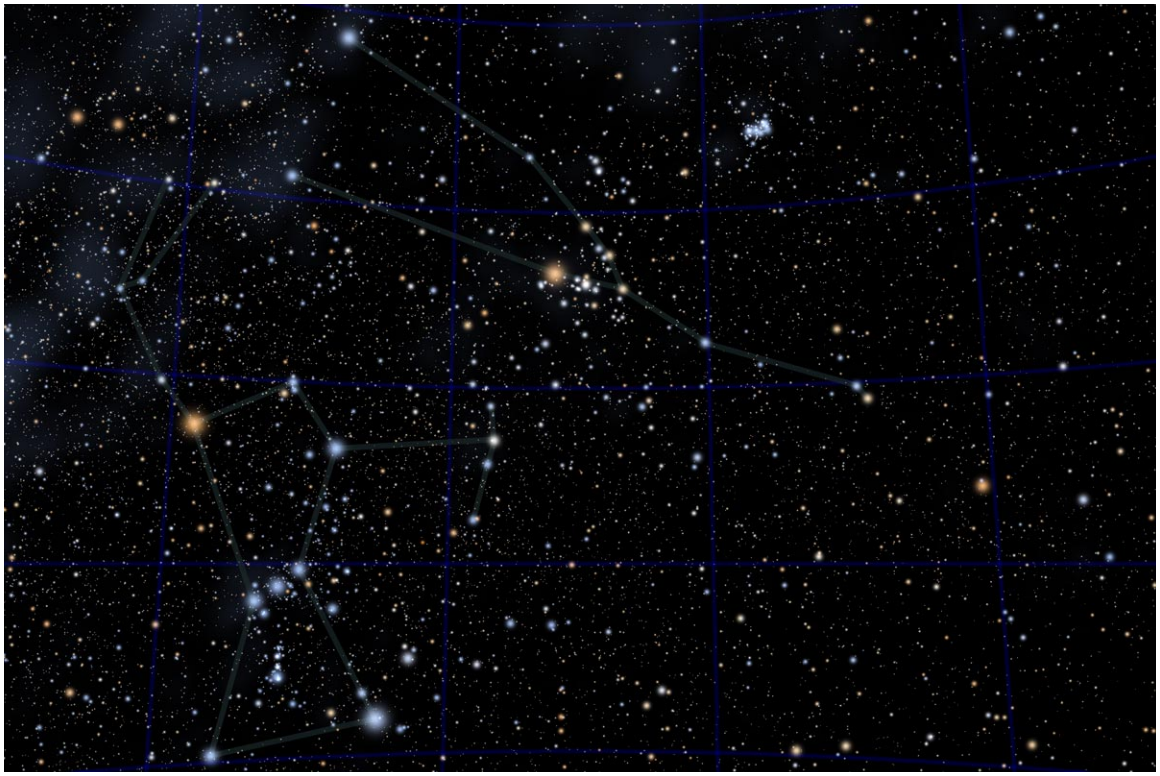

Figure 13. A star map showing Taurus and Orion using the color and brightness relations described in this paper.

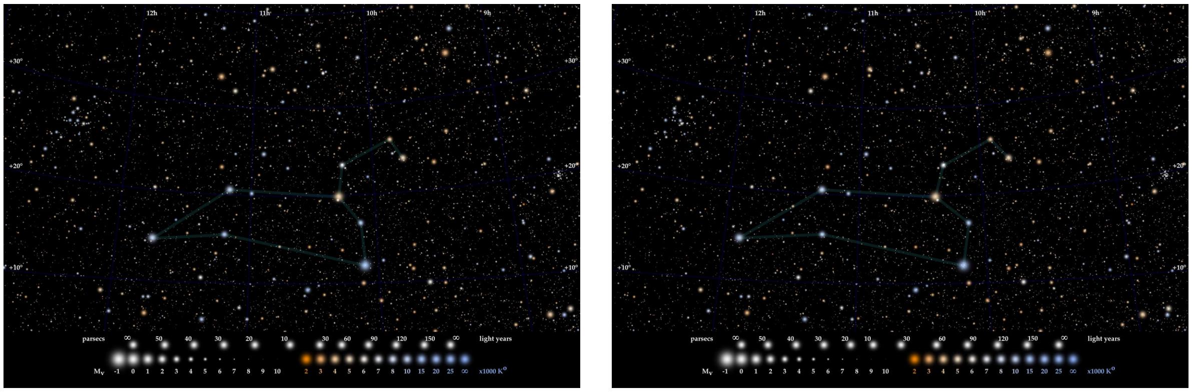

Figure 14. A stereo pair of star images showing the constellation Leo "as God might see it" (with full color vision and interocular distance of 1.6 light years). Free-view them or fuse the images using a stereo viewer. If free-viewed with crossed eyes, the distances will be inverted: the bright constellation will appear to be behind a foreground curtain of faint stars.