Night Rendering

Henrik Wann Jensen, Stanford University
Simon Premože, University of Utah
Peter Shirley, University of Utah
William B. Thompson, University of Utah
James A. Ferwerda, Cornell University
Michael M. Stark, University of Utah

Abstract

The issues of realistically rendering naturally illuminated scenes at night are examined. This requires accurate models for moonlight, night skylight, and starlight. In addition, several issues in tone reproduction are discussed: eliminatiing high frequency information invisible to scotopic (night vision) observers; representing the flare lines around stars; determining the dominant hue for the displayed image. The lighting and tone reproduction are shown on a variety of models.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism— [I.6.3]: Simulation and Modeling—Applications

Keywords: realistic image synthesis, modeling of natural phenomena, tone reproduction

1 Introduction

Most computer graphics images represent scenes with illumination at daylight levels. Fewer images have been created for twilight scenes or nighttime scenes. Artists, however, have developed many techniques for representing night scenes in images viewed under daylight conditions, such as the painting shown in Figure 1. The ability to render night scenes accurately would be useful for many applications including film, flight and driving simulation, games, and planetarium shows. In addition, there are many phenomena only visible to the dark adapted eye that are worth rendering for their intrinsic beauty. In this paper we discuss the basic issues of creating such nighttime images. We create images of naturally illuminated scenes, so issues related to artificial light sources are not considered. To create renderings of night scenes, two basic issues arise that differ from daylight rendering:

- What are the spectral and intensity characteristics of illumination at night?
- How do we tone-map images viewed in day level conditions so that they "look" like night?

Illumination computations

To create realistic images of night scenes we must model the characteristics of nighttime illumination sources, both in how much light they contribute to the scene, and what their direct appearance in the sky is:

- **The Moon**: Light received directly from the Moon, and moonlight scattered by the atmosphere, account for most of the available light at night. The appearance of the Moon itself must also be modeled accurately because of viewers' familiarity with its appearance.
- **The Sun**: The sunlight scattered around the edge of the Earth makes a visible contribution at night. During "astronomical" twilight the sky is still noticeably bright. This is especially important at latitudes more than 48° N or S where astronomical twilight lasts all night in midsummer.

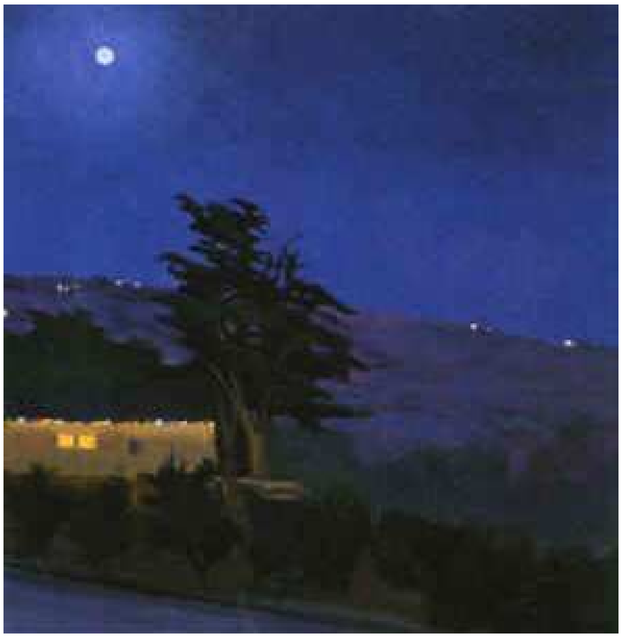

Figure 1: *A painting of a night scene. Most light comes from the Moon. Note the blue shift, and that loss of detail occurs only inside edges; the edges themselves are not blurred. (Oil, Burtt, 1990)*

- **The planets and stars**: Although the light received from the planets and stars is important as an illumination source only on moonless nights, their appearance is important for night scenes.
- **Zodiacal light**: The Earth is embedded in a dust cloud which scatters sunlight toward the Earth. This light changes the appearance and the illumination of the night sky.
- **Airglow**: The atmosphere has an intrinsic emission of visible light due to photochemical luminescence from atoms and molecules in the ionosphere. It accounts for one sixth of the light in the moonless night sky.

Several authors have examined similar issues of appearance and illumination for the daylight sky [8, 20, 32, 35, 43, 33]. To our knowledge, this is the first computer graphics paper that examines physically-based simulation of the nighttime sky. We restrict ourselves to natural lighting, and we include all significant natural sources of illumination except for aurora effects (northern lights).

Tone mapping

To display realistic images of night scenes, we must apply *tone mapping*. This is the process of displaying an image to a viewer adapted to the display environment that suggests the experience of an observer adapted to the level of illumination depicted in the scene. For our purposes this usually means displaying an image that "looks like" night to an observer that is not dark adapted. The tone mapping of night images requires us to deal with three perceptual dimensions for mesopic (between day and night vision) and scotopic (night vision) conditions (Figure 2):

- **Intensity**: How scotopic luminances are mapped to image luminances.
- **Spatial detail**: How the glare effects and loss-of-detail at scotopic levels is applied in the displayed image.
- **Hue**: How the hue of displayed image is chosen to suggest darkness in the scene.

The psychophysical approach of making displayed synthetic images have certain correct objective characteristics was introduced by Upstill [49]. The mapping of intensity has been dealt with by brightness matching [48], contrast detection threshold mapping [10, 23, 49, 52], and a combination of the two [34]. Our paper uses existing methods in intensity mapping. The loss of spatial detail has previously been handled by simple filtering to reduce high-frequency detail at scotopic levels [10, 23, 34]. This has led to an unsatisfactory blurry appearance which we attempt to address. The glare experienced in night conditions has been simulated in computer graphics [29, 42]. We use this work, and show how it should be applied for stars based on observations of stellar point-spread functions from the astronomy literature. Color shifting toward blue to suggest dark scenes is a well-known practice in film and painting, and has been partially automated by Upstill [49]. We examine the magnitude and underlying reasons for this practice and attempt to automate it. Unfortunately, this places us in the awkward position of combining empirical practices and known constraints from psychophysics. Fortunately, we can do these seemingly at-odds tasks in orthogonal perceptual dimensions. We also discuss the sensitivity of image appearance to display conditions such as background luminance and matte intensity.

The remainder of the paper is divided into two initial sections on light transport and physics of the nighttime sky, and how to perform tone reproduction from computed radiances, and is followed by example images and discussion. Our basic approach is to strive for as much physical and psychophysical accuracy as possible. This is true for the appearance of the Moon, the sky, and stars, for the amount of light they illuminate objects with, and for the tone mapping. Thus we use available height field data for the Moon topography and albedo as well as stellar position. This might be considered excessive for many applications, but it ensures that the amount of light coming from the Moon is accurate, and allows viewers with stellar navigation skills to avoid disorientation. More importantly, using real data captured with extraterrestrial measurement allows us to avoid the multidimensional parameter-tuning that has proven extremely time-consuming in production environments. However, the number of effects occurring in this framework is enormous, and we do not model some phenomena that do contribute to appearance, and these are specified in the appropriate sections. We close with results for a variety of scenes.

2 Night illumination

Significant natural illumination at night comes primarily from the Moon, the Sun (indirectly), starlight, zodiacal light and airglow. In addition a minor contribution is coming from diffuse galactic light and cosmic light. This is illustrated in Figure 3. These components of the light of the night sky can be treated separately as they are only indirectly related. Each component has two important properties for our purpose: the direct appearance, and the action as an illumination source.

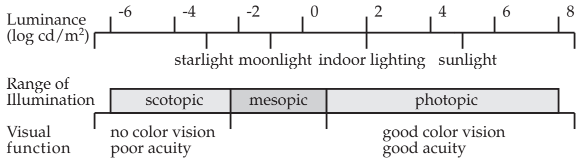

Figure 2: *The range of luminances in the natural environment and associated visual parameters. After Hood (1986)*

Component	Irradiance $[W/m^2]$
Sunlight	$1.3 \cdot 10^3$
Full moon	$2.1 \cdot 10^{-3}$
Zodiacal light	$1.2 \cdot 10^{-7}$
Integrated starlight	$3.0 \cdot 10^{-8}$
Airglow	$5.1 \cdot 10^{-8}$
Diffuse galactic light	$9.1 \cdot 10^{-9}$
Cosmic light	$9.1 \cdot 10^{-10}$

Figure 3: Typical values for sources of natural illumination at night.

The atmosphere also plays an important role in the appearance of the night sky. It scatters and absorbs light and is responsible for a significant amount of indirect illumination. An accurate representation of the atmosphere and the physics of the atmosphere is therefore necessary to accurately depict night illumination.

In this section we describe the components of the night sky and how we integrate these in our simulation.

2.1 Moonlight

To accurately render images under moonlight, we take a direct modeling approach from an accurate model of the Moon position, the measured data of the lunar topography and albedo [30]. This ensures the Moon's appearance is correct even in the presence of oscillation-like movements called *optical librations*. The Moon keeps the same face turned towards the Earth, but this does not mean that we only see half of the lunar surface. From the Earth about 59% of the lunar surface is sometimes visible while 41% of the surface is permanently invisible. Optical librations in latitude, and longitude and diurnal librations expose an extra 18% of the Moon's surface at different times. These oscillations are caused by the eccentricity of the lunar orbit (librations in longitude), deviations in inclination of the lunar rotation axis from the orbital plane (librations in latitude) and displacements in the diurnal parallax.

Position of the Moon

To compute the positions of the Sun and the Moon, we used the formulas given in Meeus [28]. For the Sun, the principal perturbations from the major planets are included, making the accuracy about ten seconds of arc in solar longitude. Meeus employs 125 terms in his formula, and it is accurate to about 15 seconds of arc. The positions of the Sun and the Moon are computed in ecliptic coordinates, a coordinate system based on the plane of the Earth's orbit, and thus requires an extra step in the transformation to be converted to altitude and azimuth. Also there are other corrections which must be applied, most notably diurnal parallax can affect the apparent position of the Moon by as much as a degree.

Illumination from the Moon

Most of the moonlight visible to us is really sunlight scattered from the surface of the Moon in all directions. The sun is approximately 500000 brighter than the moon. The Moon also gets illuminated by sunlight indirectly via the Earth.

The Moon reflects less than 10% of incident sunlight and is therefore considered a poor reflector. The rest of the sunlight gets absorbed by the lunar surface, converted into heat and re-emitted at different wavelengths outside the visible range. In addition to visible sunlight, the lunar surface is continuously exposed to high energy UV and X-rays as well as corpuscular radiation from the Sun. Recombination of atoms ionized by this high energy particles give rise to *luminescent emission* at longer wavelengths, which can penetrate our atmosphere and become visible on the ground [21]. Depending on the solar cycle and geomagnetic planetary index, the brightness of the Moon varies 10%-20%. In our simulation we do not take into account luminescent emission.

Using the Lommel-Seeliger law [51] the irradiance E_m from the Moon at phase angle α and a distance d ($r_m \ll d$) can be expressed as:

$$E_m(\alpha, d) = \frac{2}{3} C \frac{r_m^2}{d^2} (E_{s,m} + E_{e,m}) \left\{1 - \sin\left(\frac{\alpha}{2}\right) \tan\left(\frac{\alpha}{2}\right) \log\left(\cot\left(\frac{\alpha}{4}\right)\right)\right\} \quad (1)$$

where r_m is the radius of the Moon, $E_{s,m}$ and $E_{e,m}$ is the irradiance from the Sun and the Earth at the surface of the Moon. The normalizing constant C can be approximated by the average albedo of the Moon ($C = 0.072$).

The lunar surface consists of a layer of a porous pulverized material composed of particles larger than wavelengths of visible light. As a consequence and in accordance with the Mie's theory [50], the spectrum of the Moon is distinctly *redder* than the Sun's spectrum. The exact spectral composition of the light of the Moon varies with the phase (being the bluest at the full Moon). Also, previously mentioned luminescent emissions also contribute to the redder spectrum of the moonlight. In addition, light scattered from the lunar surface is polarized, but we omit polarization in our model.

Appearance of the Moon

The large-scale topography of the Moon is visible from the Earth's surface, so we simulate the appearance of the Moon in the classic method of rendering illuminated geometry and texture. The Moon is considered a poor reflector of visible light. On average only 7.2% of the light is reflected although the reflectance varies from 5-6% for crater floors to almost 18% in some craters [21]. The albedo is approximately twice as large for red (longer wavelengths) light than for blue (shorter wavelengths) light.

Given the lunar latitude β, the lunar longitude λ, the lunar phase angle α and the angle ϕ between the incoming and outgoing light, the BRDF, f, of the Moon can be approximated with [12]

$$f(\theta_i, \theta_r, \phi) = \frac{2}{3\pi} K(\beta, \lambda) B(\alpha, g) S(\phi) \frac{1}{1 + \cos\theta_r / \cos\theta_i} \quad (2)$$

where θ_r and θ_i is the angle between the reflected resp. incident light and the surface normal, $K(\beta, \lambda)$ is the lunar albedo, $B(\alpha, g)$ is a retrodirective function and S is the scattering law for individual objects.

The retrodirective function $B(\alpha, g)$ is given by

$$B(\alpha, g) = \begin{cases} 2 - \frac{\tan(\alpha)}{2g} \left(1 - e^{\frac{-g}{\tan(\alpha)}}\right) \left(3 - e^{\frac{-g}{\tan(\alpha)}}\right), \alpha < \pi/2 \\ 1, \alpha \geq \pi/2 \end{cases} \quad (3)$$

Figure 4: *The Moon rendered at different times of the month and day and under different weather conditions.*

where g is a surface density parameter which determines the sharpness of the peak at the full Moon. If ρ is the fraction of volume occupied by solid matter [12] then

$$g = k\rho^{2/3} \quad (4)$$

where k ($k \approx 2$) is a dimensionless number. Most often the lunar surface appears plausible with $g = 0.6$, although values between 0.4 (for rays) and 0.8 (for craters) can be used. The scattering law S for individual objects is given by [12]:

$$\rho_o(\phi) = \frac{\sin(|\phi|) + (\pi - |\phi|)\cos|\phi|}{\pi} + t\left(1 - \cos(|\phi|)\right)^2 \quad (5)$$

where t introduces small amount of forward scattering that arises from large particles that cause diffraction [37]. $t = 0.1$ fits Rougier's measurements [15] of the light from the Moon well.

This function has been found to give a good fit to measured data, and the complete Hapke-Lommel-Seeliger model provides good approximation to the real appearance of the Moon even though there are several discrepancies between the model and measured photometric data. Two places it falls short are opposition brightening and earthshine.

When the Moon is at opposition (ie. opposite the Sun), it is substantially brighter ($10\% - 20\%$) than one would expect due to increased area being illuminated. The mechanisms responsible for the lunar opposition effect is still being debated, but *shadow hiding* is the dominant mechanism for the opposition effect. Shadow hiding is caused by the viewer's line-of-sight and the direction of the sunlight being coincident at the antisolar point, therefore effectively removing shadows from the view. Since the shadows cannot be seen the landscape appears much brighter. We do not model opposition effect explicitly, although it affects the appearance of the Moon quite significantly. Fully capturing this effect would require a more detailed representation of the geometry than what is currently available.

When the Moon is a thin crescent the faint light on the dark side of the Moon is *earthshine*. The earthshine that we measure on the dark side of the Moon, naturally, depends strongly on the phase of the Earth. When the Earth is new (at full Moon), it hardly illuminates the Moon at all and the earthshine is very dim. On the other hand, when the Earth is full (at new Moon), it casts the greatest amount of light on the Moon and the earthshine is relatively bright and easily observed by the naked eye. If we could observe the Earth from the Moon, the full Earth would appear about 100 times brighter than the full Moon appears to us, because of the greater solid angle that the Earth subtends and much higher albedo (Earth's albedo is approximately 30%). Earthshine fades in a day or two because the amount of light available from the Earth decreases as the Moon moves around the Earth in addition to loss of the opposition brightening. We model earthshine explicitly by including the Earth as a second light source for the Moon surface. The intensity of the Earth is computed based on the Earth phase (opposite of Moon phase), the position of the Sun and the albedo of the Earth [51].

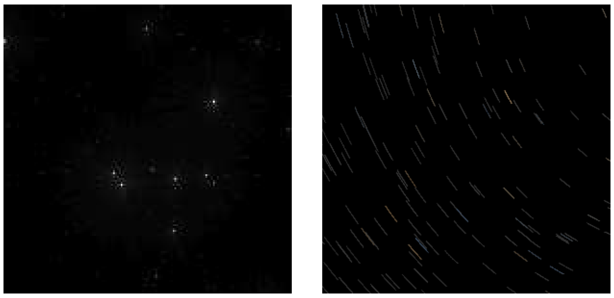

Figure 5: *Left: close-up of rendered stars. Right: close-up of time-lapsed rendering of stars.*

2.2 Starlight

Stars are obviously important visual features in the sky. While one could create plausible stars procedurally, there are many informed viewers who would know that such a procedural image was wrong. Instead we use actual star positions, brightnesses, and colors. Less obvious is that starlight is also scattered by the atmosphere, so that the sky between the stars is not black, even on moonless nights. We simulate the this effect in addition to the appearance of stars.

Position of stars

For the stars, we used the Yale Bright Star Catalog [14] of most ($\approx$ 9000) stars up to magnitude +6.5 as our database. This database contains all stars visible to the naked eye. The stellar positions are given in spherical equatorial coordinates of right ascension and declination, referred to the standard epoch J2000.0 [11].

These coordinates are based on the projection of the Earth's equator of January 1, 2000 1GMT, projected onto the celestial sphere. To render the stars, we need the altitude above the horizon and the azimuth (the angle measured west from due south) of the observer's horizon at a particular time. The conversion is a two-step process. The first is to correct for precession and nutation, the latter being a collection of short-period effects due to other celestial objects [11]. The second step is to convert from the corrected equatorial coordinates to the horizon coordinates of altitude and azimuth. The combined transformation is expressed as a single rotation matrix which can be directly computed as a function of the time and position on the Earth. We store the stellar database in rectangular coordinates, then transform each star by the matrix and convert back to spherical coordinates to obtain altitude and azimuth. Stellar proper motions, the slow apparent drifting of the stars due to their individual motion through space, are not included in our model.

Color and brightness of stars

Apparent star brightnesses are described as stellar magnitudes. The visual magnitude m_v of a star is defined as [22]:

$$m_v = -(19 + 2.5 \log_{10}(E_s)) \tag{6}$$

where E_s is the irradiance at the earth. Given visual magnitude the irradiance is:

$$E_s = 10^{-m_v - 19} \ \frac{W}{m^2} \tag{7}$$

For the Sun $m_v \approx -26.7$, for the full Moon $m_v \approx -12.2$, for Sirius (the brightest star) $m_v \approx -1.6$. The naked eye can see stars with a stellar magnitude up to approximately 6.

The color of the star is not directly available as a measured spectrum. Instead astronomers have established a standard series of measurements in particular wavebands. A widely used *UBV system* introduced by Johnson [16] isolates bands of the spectrum in the blue intensity B ($\lambda = 440nm$, $\Delta\lambda = 100nm$), yellow-green intensity V ($\lambda = 550nm$, $\Delta\lambda = 90nm$) and ultra-violet intensity U ($\lambda = 350nm$, $\Delta\lambda = 70nm$). The difference $B - V$ is called a *color index* of a star, and it is a numerical measurement of the color. A negative value of $B - V$ indicates more bluish color while a positive value indicates redder stars. UBV is not directly useful for rendering purposes. However, we can use the color index to estimate a star's temperature. [41]:

$$T_{eff} = \frac{7000K}{B - V + 0.56} \tag{8}$$

Stellar spectra are very similar to spectra of black body radiators, but there are some differences. One occurs in the *Balmer continuum* at wavelengths less than 364.6 nm largely due to absorption by hydrogen.

To compute spectral irradiance from a star given T_{eff} we first compute a non spectral irradiance value from the stellar magnitude using equation 7. We then use the computed value to scale a normalized spectrum based on Planck's radiation law for black body radiators [40]. This gives us spectral irradiance.

Illumination from stars

Even though many stars are not visible to the naked eye there is a contribution from all stars when added together. This is can be seen on clear nights as the Milky Way. Integrated starlight depends on galactic coordinates — it is brighter in the direction of the galactic plane. We currently assume that the illumination from the stars is constant ($3 \cdot 10^{-8} \frac{W}{m^2}$) over the sky. The appearance of a Milky Way is simulated using a direct visualization of the brightest stars as described in the following section.

Appearance of stars

Stars are very small and it is not practical to use ray tracing to render stars. Instead we use an image based approach in which a separate star image is generated. In addition a spectral alpha image is generated. The alpha map records for every pixel the visibility of objects beyond the atmosphere. It is generated by the renderer as a secondary image. Every time a ray from the camera exits the atmosphere and enters the universe the integrated optical depth is stored in the alpha image. The star image is multiplied with the alpha image and added to the rendered image to produce the final image. The use of the alpha image ensures that the intensity of the stars is correctly reduced due scattering and absorption in the atmosphere.

2.3 Zodiacal light

The Earth co-orbits with a ring of dust around the Sun. Sunlight scatters from this dust and can be seen from the Earth as *zodiacal light* [7, 36]. This light first manifests itself in early evening as a diffuse wedge of light in the southwestern horizon and gradually broadens with time. During the course of the night the zodiacal light becomes wider and more upright, although its position relative to the stars shifts only slightly [5]. On dark nights another source of faint light can be observed. A consequence of preferential backscattering near the anti solar point, the *gegenshine* is yet another faint source of light that is not fixed in the sky [39]. It is part of the zodiacal light. In the northern hemisphere, the zodiacal light is easiest to see in September and October just before sunrise from a very dark location.

The structure of the interplanetary dust is not well-described and to simulate zodiacal light we use a table with measured values [38]. Whenever a ray exits the atmosphere we convert the direction of the ray to ecliptic polar coordinates and perform a bilinear lookup in the table. This works since zodiacal light changes slowly with direction and has very little seasonal variation.

2.4 Airglow

Airglow is faint light that is continuously emitted by the entire upper atmosphere with a main concentration at around 110 km elevation. The upper atmosphere of the Earth is continually being bombarded by high energy particles, mainly from the Sun. These particles ionize atoms and molecules or dissociate molecules and in turn cause them to emit light in particular spectral lines (at discrete wavelengths). As the emissions come primarily from Na and O atoms as well as molecular nitrogen and oxygen the emission lines are easily recognizable. The majority of the airglow emissions occur at $557.7nm$ $(O-I)$, $630nm$ $(O-I)$ and a 589.0nm - 589.6nm doublet $(Na-I)$. Airglow is the principal source of light in the night sky on moonless nights. Airglow has significant diurnal variations. $630nm$ emission is a maximum at the beginning of the night, but it decreases rapidly and levels off for the remainder of the night [27]. All three emissions show significant seasonal changes in both monthly average intensities as well as diurnal variations.

Airglow is integrated into the simulation by adding an active layer to the atmosphere that contributes with spectral in-scattered radiance.

2.5 Diffuse galactic light and cosmic light

Diffuse galactic light and cosmic light are the last components of the night sky that we include in our simulation. These are very faint (see Figure 3) and modeled as a constant term $(1 \cdot 10^{-8} W/m^2)$ that is added when a ray exits the atmosphere.

2.6 The atmosphere modeling

Molecules and aerosols (dust, water drops and other similar sized particles) are the two main constituents of the atmosphere that affect light transport. As light travels through the atmosphere it can be scattered by molecules (Rayleigh scattering) or by aerosols (Mie scattering). The probability that a scattering event occurs is proportional to the local density of molecules and aerosols and the optical path length of the light. The two types of scattering are very different: Rayleigh scattering is strongly dependent on the wavelength of the light and it scatters almost diffusely, whereas aerosol scattering is mostly independent of the wavelength but with a strong peak in the forward direction of the scattered light.

We model the atmosphere using a spherical model similar to Nishita et al. [33] and we use the same phase functions they did to approximate the scattering of light. To simulate light transport with multiple scattering we use distribution ray tracing combined with ray marching. A ray traversing the atmosphere uses ray marching to integrate the optical depth and it samples the in-scattered indirect radiance at random positions in addition to the direct illumination. Each ray also keeps track of the visibility of the background, and all rays emanating from the camera save this information in the alpha image. This framework is fairly efficient since the atmosphere is optically thin, and it is very flexible and allows us to integrate other components in the atmosphere such as clouds and airglow.

We model clouds procedurally using an approach similar to [9] — instead of points we use a turbulence function to control the placement of the clouds. When a ray traverses a cloud medium it spawns secondary rays to sample the in-scattered indirect radiance as well as direct illumination. This is equivalent to the atmosphere but since clouds have a higher density the number of scattering events will be higher. The clouds also keep track of the visibility of the background for the alpha image; this enables partial visibility of stars through clouds. For clouds we use the Henyey-Greenstein phase-function [13] with strong forward scattering.

3 Tone mapping

Tone reproduction is usually viewed as the process of mapping computed image luminances into the displayable range of a particular output device. Existing tone mapping operators depend on adaptation in the human vision system to produce displayed images in which apparent contrast and/or brightness is close to that of the desired target image. This has the effect of preserving as much as possible of the visual structure in the image.

While people are not able to accurately judge absolute luminance intensities under normal viewing conditions, they can easily tell that absolute luminance at night is far below that present in daylight. At the same time, the low illumination of night scenes limits the visual structure apparent even to a well adapted observer. These two effects have significant impacts on tone mapping for night scenes.

While the loss of contrast and frequency sensitivity at scotopic levels has been addressed by Ward et al. [23] and Ferwerda et al. [10], the resulting images have two obvious subjective shortcoming when compared to good night film shots and paintings. The first is that they seem too blurry. The second is that they have a poor hue. Unfortunately, the psychology literature on these subjects deals with visibility thresholds and does not yet have quantitative explanations for suprathreshold appearance at scotopic levels. However, we do have the highly effective empirical practice from artistic fields. Our basic strategy is to apply psychophysics where we have data (luminance mapping and glare), and to apply empirical techniques where we do not have data (hue shift, loss of high frequency detail). We make an effort to not undermine the psychophysical manipulations when we apply the empirical manipulations.

3.1 Hue mapping

In many films, the impression of night scenes are implied by using a blue filter over the lens with a wide aperture [25]. Computer animation practitioners also tend to give a cooler palette for dark scenes than light scenes (e.g., Plate 13.29 in Apodaca et al. [2]). Painters also use a blue shift for night scenes [1]. This is somewhat surprising given that moonlight is warmer (redder) than sunlight; moonlight is simply sunlight reflected from a surface with an albedo larger at longer wavelengths.

While it is perhaps not obvious whether night scenes really "look" blue or whether it is just an artistic convention, we believe that the blue-shift is a real perceptual effect. The fact that rods contribute to color vision at mesopic levels is well-known as "rod

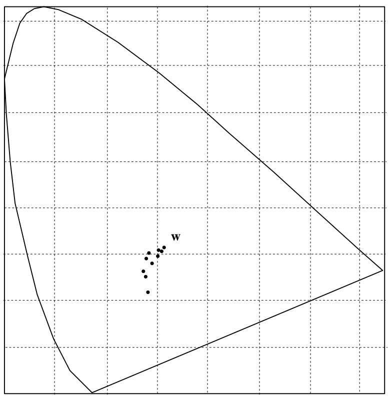

Figure 6: *The average chromaticity of "dark" sections of ten realist night paintings. The point 'w' is (0.33,0.33).*

intrusion", and is why much color research uses two degree targets to limit contribution to foveal cone cells. Trezona summarized the knowledge of rod intrusion and argues that a tetrachromatic system is needed for mesopic color perception [46]. More relevant for our scotopic renderings, in an earlier paper Trezona discussed several reasons to believe that rods contribute to the blue perceptual mechanism, and speculated that the rods may share underlying neural channels with the short-wavelength cones [45]. Upstill applied this in tone mapping by adding a partial rod response to the blue/yellow channel in his opponent processing system [49], although he was unable to set parameters precisely due to lack of specific data. However, there has not been enough quantitative study of this blue-rod effect to calibrate how we set hue in our tone-mapping. For this reason we use an empirical approach.

To gain some knowledge of artistic practice on hue-shifts for images of night scenes, we took ten realist paintings of night scenes[1], cropping off a region of pixels that appeared scotopic, applying inverse gamma to the image, and computing the average chromaticity of the pixels. We found that the average chromaticity was approximately $(x, y) = (0.29, 0.28)$ with a standard deviation of $(\sigma_x, \sigma_y) = (0.03, 0.03)$. The individual averages for the ten paintings is shown in Figure 6. Note that all of the paintings are southwest of the white point 'w'. Because we lack good statistics, and we do not have reliable measurements of the original paintings, the only conclusion to be drawn from this is that the paintings really do have a blue hue in the traditional trichromatic color science sense. Rather than try to gather more data with the same flaws, we derive a free parameter to account for this blue shift.

We assume that we have some desired dominant hue (x_b, y_b) that we would like to shift our scotopic pixels toward. We have chosen $(x_b, y_b) = (0.25, 0.25)$ as this seems to be where the bluest paintings lie and we wished to allow for the uncertainties of the printing process. As done by Upstill [49] and by Ferwerda et al. [10], we assume a scalar measure of rod saturation s. When $s = 0$ the rods are fully active and the viewer is in a fully scotopic state. When $s = 1$ the rods are fully saturated (are receiving too much light) and the viewer is in a fully photopic state. For s between 0 and 1, the viewer is in a mesopic state: 0.01 cd/m^2 to about 4.0 cd/m^2 for our purposes (Figure 2).

To accomplish the basic conversion from an *XYZV* image, we first compute $s \in [0, 1]$. Then we add a blue shift proportional to $1 - s$ so there is no shift for photopic viewers and a partial shift for mesopic viewers. Since there is not yet an accepted mesopic ($0 < s < 1$) photometry system, we compute s as a smooth transition on a log scale as indicated in the data that is available [47]. The overall computation is:

$$
\begin{aligned}
s &= \begin{cases} 0 & \text{if } \log_{10} Y < -2, \\ 3\left(\frac{\log_{10} Y + 2}{2.6}\right)^2 - 2\left(\frac{\log_{10} Y + 2}{2.6}\right)^3 & \text{if } -2 < \log_{10} Y < 0.6, \\ 1 & \text{otherwise.} \end{cases} \\
W &= X + Y + Z \\
x &= X/W \\
y &= Y/W \\
x &= (1 - s)x_b + sx \\
y &= (1 - s)y_b + sy \\
Y &= 0.4468(1 - s)V + sY \\
X &= \frac{xY}{y} \\
Z &= \frac{X}{x} - X - Y
\end{aligned}
$$

Here V is the scotopic luminance. After this process the scotopic pixels will be blue shifted, the mesopic pixels will be partially blue shifted. Also, the Y channel will store scaled scotopic luminance for scotopic pixels, photopic luminance for photopic channels, and a combination for mesopic pixels. The scaling factor 0.4468 is the ratio of Y to V for a uniform white field.

3.2 Loss of detail

The spatial density of rods is much lower than that of the foveal cones. Thus there is a loss-of-detail at scotopic levels. This is why it is impossible to read most text under moonlight. To date this effect has been simulated with low-pass filtering [10]. Subjectively, however, loss of detail in the visual field does not convey a strong sense of blurred edges. To see this, it is enough to hold your spread fingers at the side of your head and view them with peripheral vision. While the fingers appear less sharp than when viewed with foveal vision, the appearance of blurring is much less than would be expected given the limited detail and resolution in that portion of the visual field [24].

Painters simulate the loss of detail at night by retaining prominent edges, but dropping textures and other high frequency. We can approximate this in computer rendered images using a spatially varying blur, where the amount of blur is a function of the presence of absence of prominent edges. The actual technique used is patterned after that described by Cornsweet and Yellott [44], but bases the amount of blurring on a measure of local edge strength rather than local luminance. Usually salient edges in an image are characterized by a high spatial gradient, with a large spatial extent in one direction and a localized spatial extent in the orthogonal direction [26]. Modern edge detectors used in computer vision most often detect such patterns by convolving the image with a Gaussian point spread function (psf) and then examining the magnitude of the spatial gradient of the result [6]. Similarly, we based the amount of blur we used to simulate loss of detail on the magnitude of the gradient of a smoothed version of the physically correct luminance image. Unlike the situation with edge detection, no non-linear thresholding is applied to this value and so the width of the psf applied before differentiation is not critical. A value of 1.0 inter-pixel distance was used in the examples presented here.

For the final rendered image, we want more blur in regions where edges are *not* present. To accomplish this, we blur each point in the original image with a Gaussian psf having a standard deviation σ specified by

$$\sigma = \sigma_0(grad_{strong} - grad_{mag})/grad_{strong},$$

where $grad_{strong}$ is the smallest gradient magnitude for which no blurring should occur and σ_0 is the standard deviation of the point-spread function associated with the maximum blurring that should

[1] We gathered these images from *www.art.com* by searching for the keyword "moon" and eliminating paintings we evaluated as not realistic.

be applied. For the examples shown here, $grad_{strong}$ was chosen empirically, and σ_0 was chosen in the method of Ferwerda et al. [10]: the finest spatial grating pattern visible under moonlight (about four cycles per degree) was used to calibrate a value for the standard deviation which made the grating barely invisible.

3.3 Stellations

Bright stars have a familiar appearance of "stellations". It is not intuitively clear whether these are the result of atmospheric optics, human optics, or neural processing. However, careful examination of retinal images has shown that the stellations are entirely due to diffraction in the outer portion of the lens [31]. Because a display observer will have a smaller pupil than the scene observer, this diffraction effect will be absent when viewing a simulated image at non-scotopic levels. We use the scotopic filter presented by Spencer et al. [42] to achieve this effect.

3.4 Final tone mapping

To actually set the contrasts of the final displayed image, we can use any of the standard achromatic tone-reproduction operators. In particular, we use the simplest one: the contrast matching technique of Ward [52]. Fortunately, naturally illuminated night scenes have a relatively low dynamic range, so more complex operators are not crucial. For higher dynamic range scenes, such as those with artificial light, a higher dynamic range intensity mapping method could be used without changing the rest of the framework.

3.5 Viewing conditions

Renderings of night scenes are sensitive to viewing conditions. In this section we discuss some of the known issues that are expected to affect the perceived display.

Prints, backlit transparencies, and projected transparencies are typically viewed under very different conditions of illumination. Prints and backlit transparencies are usually viewed under general room illumination. Under these conditions the observer's state of adaptation will be determined primarily by the photograph's surround. Projected transparencies are usually viewed under dark room conditions. In this case the observer's state of adaptation will be determined primarily by the photograph itself. Jones et al. [17, 18, 19] conducted a series of experiments to determine how different photographic media and viewing conditions under which they are seen, determines which objective tone reproduction curve provides the best subjective visual experience. They found that as the dynamic range of the medium increases (print, backlit transparency, projected transparency) the preferred gamma of the tone reproduction curve increases as well, going from about 1.1 for the print to about 1.5 for the projected transparency. A similar relationship was found in studies by Bartelson and Breneman [3, 4].

In addition to contrast adjustments, we can sometimes control the surround of an image using matting. Matting potentially serves two, rather different purposes. The first is to provide a reference luminance so that the overall image looks darker than it would with a darker surround. The second and less obvious purpose is that mats, particularly when they have a 3-D structure, provide the same sense as looking through a window. This has the perceptual effect of decoupling apparent image brightness and geometry from screen brightness and geometry.

Although we do not control for the effects of matting and larger surround effects, we caution practitioners to be aware that they can be important when viewing rendered images.

4 Results

We have integrated the night simulation framework into a global illumination renderer. This renderer samples radiance at every 10nm (40 samples) and output images in XYZV format for further tone-mapping. Global illumination of diffuse geometry is done using a spectral irradiance cache.

4.1 Terrain example

Our first test scene is based on a digital elevation model (DEM) of the Hogum Cirque. This scene has approximately 200 million triangles (terrain plus instanced plants).

Figure 7 shows the various stages of tone mapping for a rendered image. The computed illuminance from the Moon is 0.1 lux and the viewing conditions are scotopic. For the tone mapping notice how the images without the blue-shift look too red (warm) and saturated. Also, notice the blurring of the details in the texture in the final image. For these images we did not include stars.

Figure 8 shows the final tone-mapped image with stars. The subtle glare effects around the brighter stars can be seen. In figure 9 we have show tone mapped versions of Hogum under different lighting conditions. Notice how the image with full Moon illumination is brighter than the other images. The computed illuminance from the Moon in these images varies from 0.2 lux (full Moon) to 0.02 lux (half Moon). The computed illuminance from the sun in the daylight image is 121,000 lux. Notice how the clouds correctly block the light from some of the stars. For comparison we also included a daylight rendering. It is not only brighter, but also much more saturated than the tone-mapped night images. The daylight image is more than five orders of magnitude brighter than the moonlight image.

Figure 10 shows a moon setting over the ridge of Hogum. Notice that the moonlight is scattered by the clouds (aureole). Figure 11 shows an example of Hogum covered in snow. This blue-shift becomes very noticeable and the brightness of the snow begins to dominate over the sky.

Each of the Hogum images with clouds and global illumination was rendered in approximately 2 hours on a quad PII-400 PC using less than 60MB of memory.

4.2 Zodiacal light and star trails

In figure 12 we have rendered a view in the direction of the rising sun at a time when the sun is still 20 degrees below the horizon and therefore constributes only very little illumination to the sky. In addition the atmosphere is clear and the moon is not visible. This exposes the zodiacal light in the direction of the ecliptic plane. It can be seen as a wedge of light rising from the horizon. Images of zodiacal light normally have long exposure times. This is also the case for this image and the reason why the startrails appear.

5 Conclusion and future work

This paper has delved into several issues of night rendering. Each of the computational methods has limitations, but we think the basic parameters of night rendering under natural illumination have been addressed. Obvious places for improvement are the inclusion of artificial lighting, planets, and dynamic adaptation effects as occur when stepping from illuminated rooms into the night. The biggest limitation to the work we have presented is the empirical adjustment of parameters for hue shift and edge thresholds in tone mapping. This can be improved once there is better psychophysical data for suprathreshold appearance — that data is not yet known. The tone mapping could also be improved by better accounting for display

Figure 8: *Final tone-mapping of Hogum.*

conditions that influence appearance of night images viewed in non-scotopic conditions.

References

[1] George. A. Agoston. *Color theory and its application in art and design*. Springer-Verlag, 1979.

[2] Anthony A. Apodaca and Larry Gritz. *Creating CGI for Motion Pictures*. Morgan Kaufman, 1999.

[3] C. J. Bartleson and E. J. Breneman. Brightness perception in complex fields. *Journal of the Optical Society of America*, 57(7):953–957, July 1967.

[4] C. J. Bartleson and E. J. Breneman. Brightness reproduction in the photographic process. *Photographic Science and Engineering*, 11(4):254–262, July-August 1967.

[5] D. E. Blackwell. The zodiacal light. *Scientific American*, 54, July 1960.

[6] J.F. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8(6):679–698, November 1986.

[7] S. F. Dermott and J. C. Liou. Detection of asteroidal dust particles from known families in near-earth orbits. In *AIP Conference Proceedings*, volume 301, pages 13–21, July 1994.

[8] Yoshinori Dobashi, Tomoyuki Nishita, Kazufumi Kaneda, and Hideo Yamashita. A fast display method of sky colour using basis functions. *The Journal of Visualization and Computer Animation*, 8(3):115–127, April – June 1997. ISSN 1049-8907.

[9] David S. Ebert, F. Kenton Musgrave, Darwyn Peachey, Ken Perlin, and Steven Worley. *Texturing and Modeling: A procedural Approach*. Academic Press, 1994.

[10] James A. Ferwerda, Sumant Pattanaik, Peter Shirley, and Donald P. Greenberg. A model of visual adaptation for realistic image synthesis. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 249–258. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.

[11] Robin M. Green, editor. *Spherical Astronomy*. Cambridge University Press, 1985.

[12] Bruce W. Hapke. A theoretical photometric function of the lunar surface. *Journal of Geophysical Research*, 68(15):4571–4586, 1963.

[13] L. G. Henyey and J. L. Greenstein. Diffuse radiation in the galaxy. *Astrophysics Journal*, 93:70–83, 1941.

[14] D. Hoffleit and W. Warren. *The Bright Star Catalogue*. Yale University Observatory, 5th edition, 1991.

[15] J. van Diggelen. Photometric properties of lunar carter floors. *Rech. Obs. Utrecht*, 14:1–114, 1959.

[16] H. L. Johnson and W. W. Morgan. Fundamental stellar photometry for standards of spectral type on the revised system of the yerkes spectral atlas. *Astrophysics Journal*, 117(313), 1953.

[17] L. A. Jones and H. R. Condit. Sunlight and skylight as determinants of photographic exposure: I. luminous density as determined by solar altitude and atmospheric conditions. *Journal of the Optical Society of America*, 38(2):123–178, February 1948.

[18] L. A. Jones and H. R. Condit. Sunlight and skylight as determinants of photographic exposure: II. scene structure, directional index, photographic efficiency of daylight and safety factors. *Journal of the Optical Society of America*, 39(2):94–135, February 1949. see page 123.

[19] L. A. Jones and C. N. Nelson. Control of photographic printing by measured characteristics of the negative. *Journal of the Optical Society of America*, 32(10):558–619, October 1942.

[20] R. Victor Klassen. Modeling the effect of the atmosphere on light. *ACM Transactions on Graphics*, 6(3):215–237, 1987.

[21] Zdeněk Kopal. *The Moon*. D. Reidel Publishing Company, Dordrecht, Holland, 1969.

[22] Kenneth R Lang. Astrophysical formulae / k.r. lang.

[23] Gregory Ward Larson, Holly Rushmeier, and Christine Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October - December 1997. ISSN 1077-2626.

[24] Jack Loomis. Personal Communication.

[25] Kris Malkiewicz. *Film Lighting: Talks With Hollywood's Cinematographers and Gaffers*. Simon & Schuster, 1992.

[26] D. Marr and E. Hildreth. Theory of edge detection. *Proceedings of the Royal Society London*, B 207:187–217, 1980.

[27] Billy M. McCormac, editor. *Aurora and Airglow*. Reinhold Publishing Corporation, 1967.

[28] Jean Meeus. *Astronomical Formulae for Calculators*. Willman-Bell, Inc., 4th edition, 1988.

[29] Eihachiro Nakamae, Kazufumi Kaneda, Takashi Okamoto, and Tomoyuki Nishita. A lighting model aiming at drive simulators. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 395–404, August 1990.

[30] Naval Research Laboratory. Clementine deep space program science experiment. http://www.nrl.navy.mil/clementine/.

[31] R. Navarro and M. A. Losada. Shape of stars and optical quality of the human eye. *Journal of the Optical Society of America (A)*, 14(2):353–359, 1997.

[32] Tomoyuki Nishita and Eihachiro Nakamae. Continuous tone representation of three-dimensional objects illuminated by sky light. In *Computer Graphics (SIGGRAPH '86 Proceedings)*, volume 20, pages 125–32, August 1986.

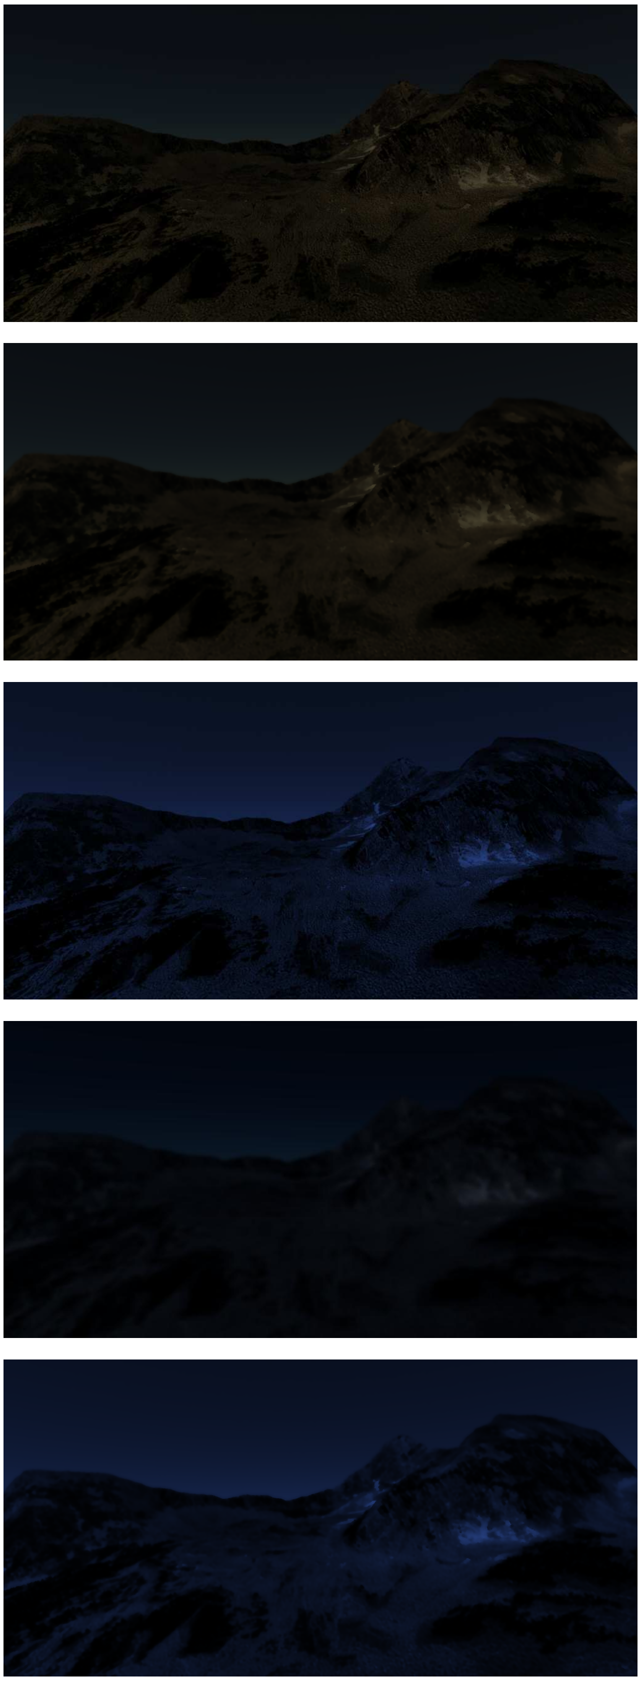

Figure 7: *Various dimensions of tone-mapping left out. Top to bottom: Ward's method without enhancements; No blue shift; No loss-of-detail; No selective loss-of-detail; Full tone reproduction.*

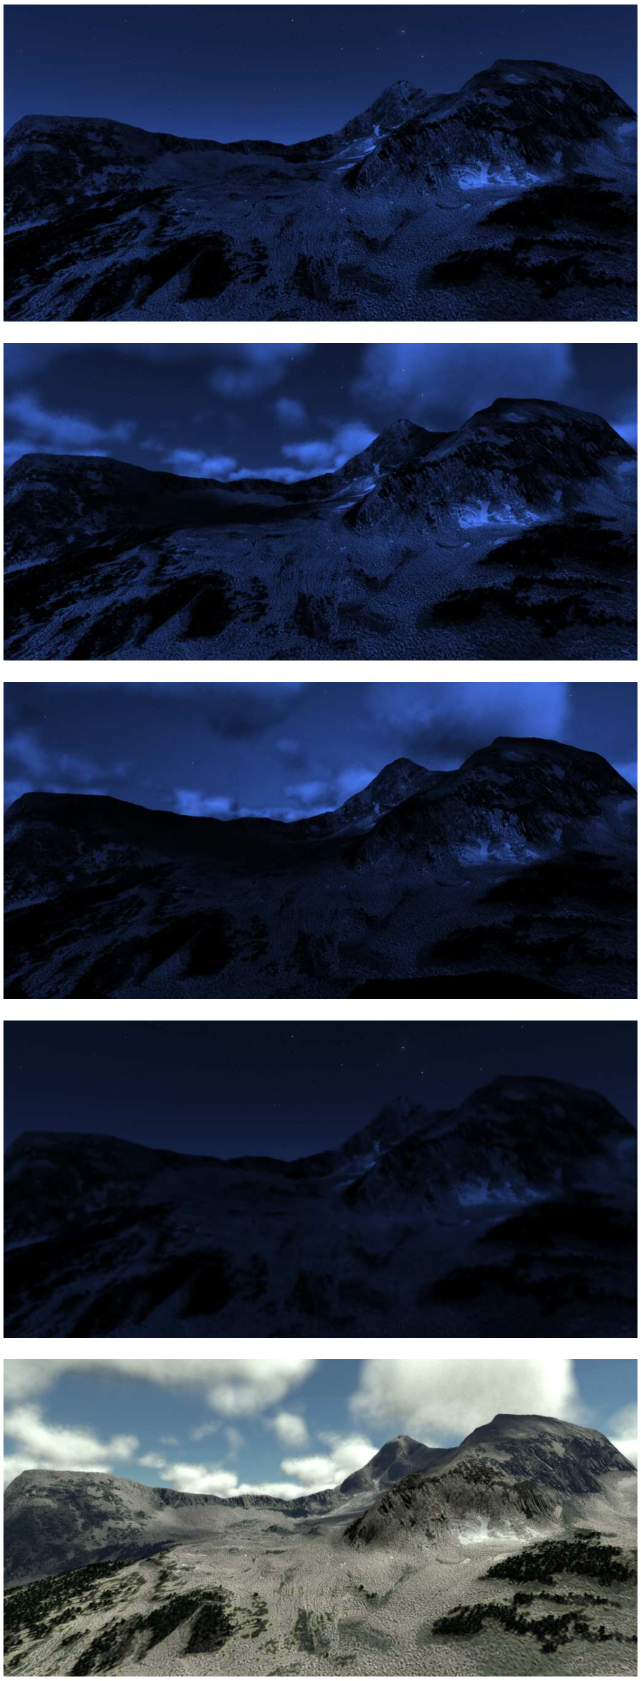

Figure 9: *Hogum mountain scene. Top to bottom: Full Moon; Full Moon with clouds; Low full Moon; Half Moon; Daylight*

Figure 10: *Moon rising above mountain ridge. Notice scattering in the thin cloud layer.*

Figure 11: *Snow covered Hogum.*

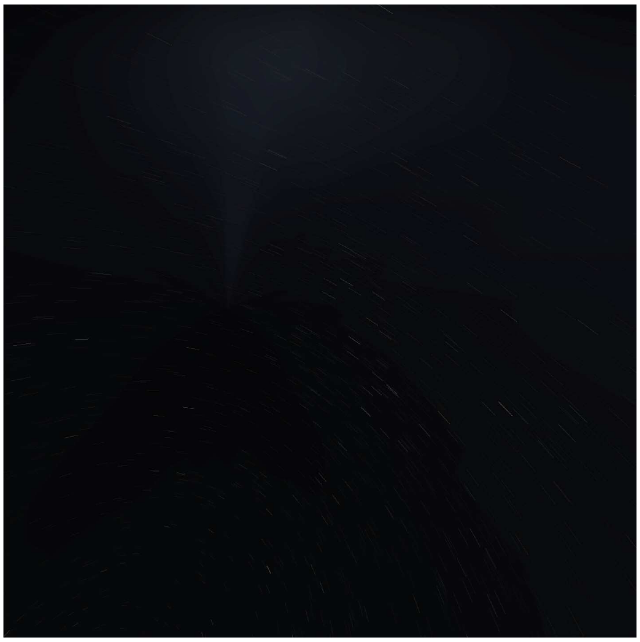

Figure 12: *Zodiacal light seen as a thin wedge of light rising from the horizon*

[33] Tomoyuki Nishita, Takao Sirai, Katsumi Tadamura, and Eihachiro Nakamae. Display of the earth taking into account atmospheric scattering. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 175–182, August 1993.

[34] Sumanta N. Pattanaik, James A. Ferwerda, Mark D. Fairchild, and Donald P. Greenberg. A multiscale model of adaptation and spatial vision for realistic image display. In Michael Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 287–298. ACM SIGGRAPH, Addison Wesley, July 1998. ISBN 0-89791-999-8.

[35] A. J. Preetham, Peter Shirley, and Brian Smits. A practical analytic model for daylight. In *SIGGRAPH*, 1999.

[36] William T. Reach, Bryan A. Franz, Thomas Kelsall, and Janet L. Weiland. Dirbe observations of the zodiacal light. In *AIP Conference Proceedings*, volume 348, pages 37–46, January 1996.

[37] N. B. Richter. The photometric properties of interplanetary matter. *Quarterly Journal of the Royal Astronomical Society*, (3):179–186, 1962.

[38] Franklin Evans Roach and Janet L. Gordon. *The Light of the Night Sky*. Geophysics and Astrophysics Monographs, V. 4). D Reidel Pub Co, June 1973.

[39] R. G. Roosen. The gegenschein and interplanetary dust outside the eart's orbit. *Icarus*, 13, 1970.

[40] Robert Siegel and John R. Howell. *Thermal Radiation Heat Transfer*. Hemisphere Publishing Corporation, 3rd edition, 1992.

[41] Robert C. Smith. *Observational Astrophysics*. Cambridge University Press, 1995.

[42] Greg Spencer, Peter Shirley, Kurt Zimmerman, and Donald Greenberg. Physically-based glare effects for digital images. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 325–334. ACM SIGGRAPH, Addison Wesley, August 1995. held in Los Angeles, California, 06-11 August 1995.

[43] Katsumi Tadamura, Eihachiro Nakamae, Kazufumi Kaneda, Masshi Baba, Hideo Yamashita, and Tomoyuki Nishita. Modeling of skylight and rendering of outdoor scenes. In R. J. Hubbold and R. Juan, editors, *Eurographics '93*, pages 189–200, Oxford, UK, 1993. Eurographics, Blackwell Publishers.

[44] T.N.Cornsweet and J.I. Yellot. Intensity dependent spatial summation. *Journal of the Optical Society of America A*, 2(10):1769–1786, 1985.

[45] P. W. Trezona. Rod participation in the 'blue' mechanism and its effect on color matching. *Vision Research*, 10:317–332, 1970.

[46] P. W. Trezona. Problems of rod participation in large field color matching. *Color Research and Application*, 20(3):206–209, 1995.

[47] P. W. Trezona. Theoretical aspects of mesopic photometry and their implication in data assessment and investigation planning. *Color Research and Application*, 23(5):264–173, 1998.

[48] Jack Tumblin and Holly E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993. also appeared as Tech. Report GIT-GVU-91-13, Graphics, Visualization & Usability Center, Coll. of Computing, Georgia Institute of Tech.

[49] Steven Upstill. *The Realistic Presentation of Synthetic Images: Image Processing in Computer Graphics*. PhD thesis, Berkeley, 1985.

[50] H.C. van de Hulst. *Light Scattering by Small Particles*. John Wiley & Sons, 1957.

[51] H.C. van de Hulst. *Multiple Light Scattering*. Academic Press, New York, NY, 1980.

[52] Greg Ward. A contrast-based scalefactor for luminance display. In Paul Heckbert, editor, *Graphics Gems IV*, pages 415–421. Academic Press, Boston, 1994.