

# Visualizing Stars and Emission Nebulae

David R. Nadeau  
Jon D. Genetti  
Steve Napear  
Bernard Pailthorpe

San Diego Supercomputer Center  
University of California, San Diego, USA

Carter Emmart  
Erik Wesselak  
Dennis Davidson

Hayden Planetarium  
American Museum of Natural History  
New York City, New York, USA

---

## Abstract

We describe the star and nebula visualization techniques used to create a 3D volumetric visualization of the Orion Nebula. The nebula's ionization layer is modeled first as a surface model, derived from infrared and visible light observations. The model is imported into a volume scene graph-based volume visualization system to simulate the nebula's emissive gases. Stars are rendered using Gaussian spots that are attenuated with distance.

---

## 1. Introduction

In 1999, the Hayden Planetarium at the American Museum of Natural History, and the San Diego Supercomputer Center at the University of California, San Diego, collaborated to create a 3D visualization of the Orion Nebula. The planetarium dome's new *digital* projection system provided a unique opportunity to use state-of-the-art computer graphics to take the audience away from Earth and examine one of our galaxy's most fascinating locales.

This section briefly explains the Orion Nebula's luminosity. Section 2 discusses using volume visualization to model the nebula. Section 3 discusses star rendering and light attenuation. Section 4 shows results.

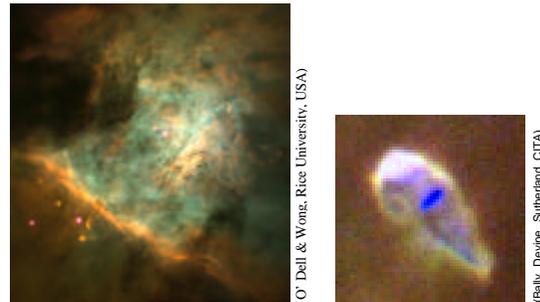
### 1.1. The Orion Nebula

Enormous clouds of dust and gas are found throughout the galaxy<sup>1,2,3</sup>. One of the closest is the Orion Nebula, shown in Figure 1.1a, which is 1500 light-years from Earth and measures several light-years across. It is visible to the human eye as a fuzzy patch in the constellation of Orion.

The galaxy contains tens of thousands of dark nebulae, so-called because the dust and gas obscure the light of stars behind them. Over time clumps of higher density gas form and grow within some of these, their gravitational attraction drawing matter from the surrounding cloud. As a clump grows, the weight of layer upon layer of gas builds up, increasing the pressure and temperature at the clump's core. The pressure continues to rise until hydrogen nuclei are packed so tightly together that they fuse, igniting a thermonuclear reaction that signals the birth of a star. We see this happening in the Orion Nebula – it is the birthplace of stars.

Hot young stars born within the nebula radiate their energy outward into the surrounding gas. High-energy photons from the stars ionize the atoms of the gas, knocking

electrons from their orbits. As these electrons collide with other electrons and slowly return to their former orbits, they emit light. It is this light we see as the nebula's eerie glow. The Orion Nebula is an example of an *emission nebula*.



**Figure 1.1.** Hubble Space Telescope images of the Orion Nebula and the HST-10 proplyd.

Since electrons can reside in atoms only in discrete energy levels, when electrons drop from outer to inner orbits they emit light at discrete wavelengths. By examining the spectra of nebulae, astronomers deduce their chemical content. Most emission nebulae are about 90% hydrogen, with the remainder helium, oxygen, nitrogen, and other elements. Ionization of these gases gives nebulae many of the colors we see in astronomical photographs.

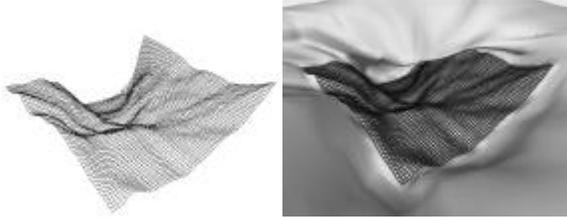
Imagery from the Hubble Space Telescope reveals dozens of stars forming within the Orion Nebula. Wrapped in cocoons of dust and gas, these protostars, or *proplyds*, often include protoplanetary disks that astronomers believe are planetary systems in the making (see Figure 1.1b).

### 1.2. The 3D Structure of the Orion Nebula

Radiation pressure from the nebula's stars pushes nearby gas away, creating cavities within the nebula's cloud. In the Orion Nebula, four centrally located hot, young stars, called

the *Trapezium*, have hollowed out the core of the nebula. This hollow core has “broken through” the portion of the cloud facing Earth, enabling us to peer inside.

Working with infrared and visible light observations from Hubble and ground-based imagery, C.R. O’Dell and Zheng Wen, of Rice University, USA, derived a 3D model of the inner surface of the hollowed out center of the nebula<sup>4</sup>. Their model shows that the region is a wrinkled, shallow “valley,” the surface of which glows from the influence of the young stars above.



**Figure 1.2.** The ionization front, derived from infrared and visible light observation (left), and the extrapolated model including the overhanging Dark Bay (right).

The ionizing effects of the trapezium’s stars penetrate a limited distance into the nebula. The glow we see is the result of a thin glowing *ionization layer* atop the valley. Dust in this surface region also reflects starlight, contributing to the total luminosity.

As in any wrinkled surface, portions of the ionization layer that face the Trapezium’s stars glow more brightly than do portions that face away. For instance, along the left side of the valley in Figure 1.2, a steep cliff faces the central stars. The cliff face is ionized and glows with a bright yellow light, seen in Hubble imagery as the *Bright Bar* feature stretching from the mid-left to the lower-right side in Figure 1.1. These lighting effects create complex rippling patterns in the nebula’s ionization layer.

Overhanging the far end of the “valley” is a dark cloud called the *Dark Bay*. The underside of this cloud is illuminated by the nebula’s stars, but the upper reaches remain dark – the Trapezium’s ionizing effects are blocked by intervening dust and gas. The Dark Bay is visible in Figure 1.1 as the dark region in the upper-left quadrant.

## 2. Visualizing the Orion Nebula

Volume visualization provides a natural approach to modeling the 3D gaseous effects of nebula. To visualize the Orion Nebula, the ionization layer model in Figure 1.2 was extrapolated outward to include the surrounding regions and overhanging Dark Bay. This model was imported as a shape in a volume modeling system. The system is based upon a *volume scene graph*<sup>5</sup>: a hierarchical organization of shapes and groups that define the content of a scene.

Volume scene graph nodes are *functions* that take a 3D location in space as an argument. When evaluated at such a location, a node function returns a color value. Most functions return values that may vary over space. Procedural

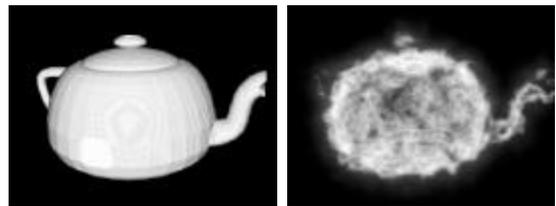
texture functions, for instance, use 3D noise, turbulence, or marble vein algorithms to compute a return value. Volume data set functions use the 3D location argument to select and interpolate voxel values from a data set read from disk.

To express the Orion Nebula’s polygonal model as a volume, scene graph nodes computed a *distance field*<sup>6,7</sup> which encodes the distance from any point in space to the nearest point on the surface. To create a soft, gaseous layer atop the surface, the field’s distances were used to vary opacity. Regions near the surface were made semi-transparent, while those inside were made opaque and those far from the surface were made transparent.



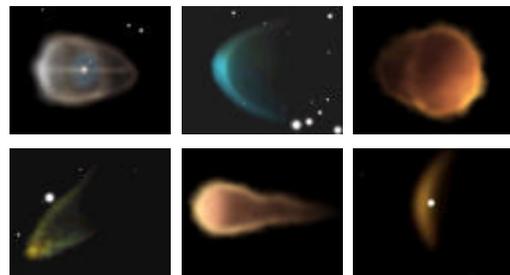
**Figure 2.1.** The ionization layer is modeled by varying opacity with distance from the polygonal model’s surface.

Using the distance field, the resulting gaseous layer is smooth and follows the terrain of the polygonal model. To give the layer a rougher, more turbulent look like that observed in real nebulae, the distance field was perturbed by 3D procedural turbulence<sup>8</sup>. Figure 2.2 illustrates the effect on the ubiquitous teapot.



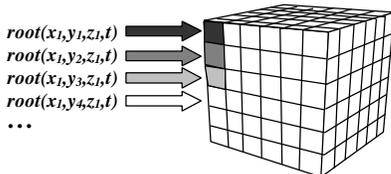
**Figure 2.2.** A distance field for a smooth surface (left), and with turbulence to deform the surface (right).

Polygonal models, volume scene graphs, distance fields, and turbulence also were used to model 85 separate proplyds and shock fronts within the nebula.



**Figure 2.3.** An assortment of proplyds and shock fronts.

*Voxelization* of the resulting volume scene graph repeatedly evaluates the graph’s functions, once for each 3D location in a grid spanning a region of interest. Each evaluation returns an RGB-alpha value that is saved into a voxel in a new volume data set. Images of the resulting data set can be created using a volume renderer.



**Figure 2.4.** Voxelization of a volume scene graph evaluates the scene graph’s functions for each 3D location on a grid.

### 2.1. Making a Nebula Glow

To render the volume, a ray-caster computes samples at regular steps along a ray. Ray samples can be composited together using the traditional formula from<sup>9</sup>:

$$c = a_1c_1 + (1 - a_1)(a_2c_2 + (1 - a_2)(...)) \quad \text{EQ.1}$$

In this compositing equation,  $c_i$  is the RGB color of the  $i^{\text{th}}$  sample along the ray, and  $\alpha_i$  is the sample’s opacity as a value from 0.0 (transparent) to 1.0 (opaque).

However, this equation presumes that an opaque sample ( $\alpha = 1.0$ ) blocks light from samples further along the ray. While this correctly models solid objects, it does not correctly model gases. The glowing gas of a nebula, for instance, emits light, but does not occlude.

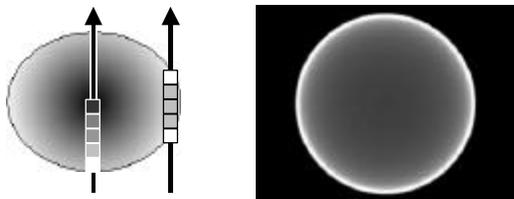
To model glowing gas effects, rendering uses a modified compositing equation that models radiative transfer<sup>1,10</sup>:

$$c = a_1c_1 + (1 - b_1)(a_2c_2 + (1 - b_2)(...)) \quad \text{EQ.2}$$

In this equation, the  $(1 - \alpha)$  terms are replaced by  $(1 - b)$  terms. The new  $b$  value encodes *opacity* (or absorptivity) while the original  $\alpha$  value encodes *emissivity*.

Dark dust in the Orion Nebula’s main cloud, and in the Dark Bay, were modeled with high opacity and low emissivity. The glowing ionization layer was modeled with low opacity and high emissivity, both of which varied with distance from the surface of the ionization front. Empty spaces in the nebula had zeroes for opacity and emissivity.

Using this approach, a shape with a high emissivity / low opacity layer surrounding a dark opaque core appears to have a halo. Figure 2.5a illustrates two rays cast toward such a shape. One ray glances across the edge of the shape, intersecting multiple samples that glow, but don’t block the glow of samples further along the ray. The emissive color of these samples sums to a bright color for the ray, creating an edge-brightening effect.



**Figure 2.5.** Edge brightening occurs when rays pass through high-emissivity, low-opacity samples at the outer layers of shapes such as this sphere.

In Figure 2.5a, a second ray strikes the shape head-on, passing through only a few emissive samples before entering the dark opaque core. When these samples are summed, the ray has a dimmer color than the first ray.

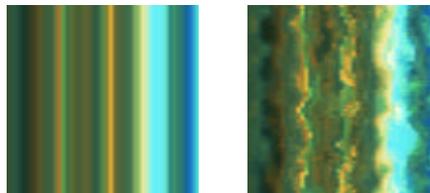
### 2.2. Shading the Nebula

Recall that the ionization front of the Orion Nebula is a wrinkled surface. To model the light and dark patterns of this surface, we could shade the volume based upon the light cast by nearby stars. At each sample along a ray, a gradient could be computed and used to approximate the surface normal. The angle between this normal and a vector to each star would provide an intensity for the sample.

Rather than compute lighting effects, we used a texture-mapping approach that reduced the computation required. We noted that the Hubble image of the Orion Nebula already encodes both color *and* intensity. What’s more, the light-and-dark texture of the nebula aligns perfectly with the wrinkles of the ionization layer – the polygonal model was generated based upon nebula images, after all.

So, to give the nebula its color and brightness, the Hubble image was projected through the volume along a vector from Earth. The image was processed first to reduce its red content, correcting for the reddening effects of intervening material between Earth and the nebula. Because stars, proplyds, and shock fronts were added later, these were painted out of the Hubble image.

Simple projection, however, would create stripes when image pixels project in a straight line through the volume. These stripes would be apparent even in the thin ionization layer. To break up these stripes, texture coordinates were jittered by procedural noise.



**Figure 2.6.** Volume cross-sections showing unjittered and jittered projection of an image through the volume.

### 3. Visualizing Stars

Using Hubble imagery, C.R. O’Dell and S.K. Wong cataloged 344 stars in the vicinity of the Orion nebula<sup>11</sup>. Their work provided the absolute magnitudes and locations of stars. Locations, however, were known only by their angular position in Earth’s sky. This places stars side-to-side within the nebula, but does not provide their elevation above the ionization layer.

Star locations were extrapolated to obtain a best-guess estimate of locations within the 3D space of the nebula. Combined with additional star data from sky surveys, a total of 884 stars were placed within the nebula region.

For all intents and purposes, a star is a point light source. An observer must get fairly close before a star resolves to a glowing sphere. So, stars could be rendered as simple dots of varying brightness.

However, in astronomical photographs the light of bright stars bleeds, creating a spot whose radius varies with star brightness. Since nebula images are invariably overexposed to make the nebula's dim glow visible, the public is used to seeing a nebula's stars as large bright spots. (Photographs also usually show cross-shaped diffraction spikes – caused by telescope mirror supports – radiating from the stars; however, we do not address mimicking these effects.)

To maintain the “realism” of the visualization, stars were rendered to mimic photographic bleeding effects. Stars were modeled as Gaussian spots whose radius and brightness varied with a star's *apparent magnitude* – the brightness of a star when viewed from a distance.

Astronomers relate a star's true brightness, or *absolute magnitude*,  $M$ , to its apparent magnitude,  $m$ , at a distance,  $d$  (in parsecs), by the formula<sup>12</sup>:

$$m = M - 5 + 5 \log(d) \quad \text{EQ.3}$$

Predictably, as distance to the star increases, the apparent magnitude decreases and the star looks dimmer.



**Figure 3.1.** A star at increasing distances.

This equation provides for the attenuation of star light. But should the light of the nebula be attenuated as well?

Consider first a point light source centered within a sphere of radius  $d$ . The amount of light passing through each square meter of the sphere each second is the total light from the source,  $L$ , divided by the sphere's surface area,  $4\pi d^2$ . The light's apparent brightness,  $b$ , is then

$$b = \frac{L}{4\pi d^2} \quad \text{EQ.4}$$

The brightness of a point light source is then inversely proportional to the square of the distance to the light source. The star attenuation equation given earlier is derived from this relationship.

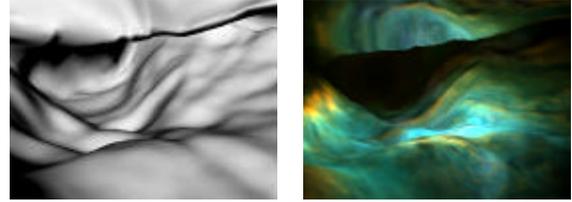
A nebula is an extended object that acts like an *area light source*, not a point light source. As above, the total amount of light received from an object decreases with the square of the distance. But the angular area of an extended object is the square of its angular extent, which varies inversely with distance. As a result, the angular “footprint” of the object on the image plane also varies with the inverse square of the distance. So, since the footprint and the total received illumination of the object both vary by the same factor, the net result is that *the surface brightness of an extended object is independent of its distance*.

This somewhat non-intuitive result means that nebulae viewed by the unaided human eye appear as bright in the night sky of Earth as they would if you were up close. This

also underscores that nebulae are, in fact, rather dim phenomena. It is only through over-exposure that we are able to see their incredible detail and colors.

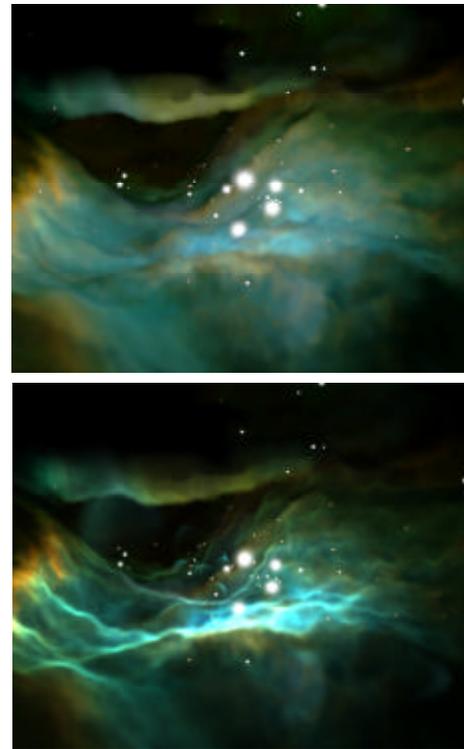
#### 4. Results

Volumetric models of the nebula and its proplyds and shock fronts were rendered, along with stars, using a multi-volume perspective ray-casting renderer<sup>13</sup>. Figure 4.1 shows the original shaded and textured polygonal model, looking down the “valley” and under the overhang of the Dark Bay.



**Figure 4.1.** The original polygonal model shaded (left) and texture mapped (right).

Figure 4.2 shows the volume rendered model. The top image shows the model rendered *without* per-voxel emissivity, just opacity (i.e., using conventional ray compositing with EQ. 1). The bottom image encodes per-voxel emissivity as well (i.e., using EQ. 2). Note the edge-brightening effects across the wrinkled surface of the ionization layer.



**Figure 4.2** The volume rendered model without (top) and with (bottom) per-voxel emissivity.

The close-up in Figure 4.3 shows some of the proplyds and shock fronts in the nebula's center.



**Figure 4.3.** A close-up showing proplyds and shock fronts in the nebula's center.

A 2 ½ minute fly-through animation of the nebula was produced for the Hayden Planetarium's daily show. The planetarium uses seven 1280x1024 video projectors to seamlessly cover the interior of its dome. Animation production, then, computed seven high-resolution images for each frame of the show. The 2 ½ minute animation required about 31,000 high-resolution images.

Animation frames were rendered using 900+ processors on the San Diego Supercomputer Center's IBM RS/6000 teraflops supercomputer. Running one multi-threaded renderer on each 8-processor node, the frames were computed during a single 12-hour period.

#### Acknowledgements

This work was funded by NASA as part of the Hayden Planetarium's *Digital Galaxy Project*. C.R. O'Dell of Rice University provided invaluable guidance on the structure of the Orion nebula, and supplied data for the ionization layer terrain, star and proplyd locations, and the de-reddened Hubble image. Greg Johnson of the San Diego Supercomputer Center (SDSC) helped in multi-threading the renderer. Mike Gannis provided expert help on this paper. The software developed is part of a larger suite of scalable volume visualization tools being developed under the NSF-funded National Partnership for Advanced Computational Infrastructure (NPACI). The scalable tools project is led by B. A. Pailthorpe and A. Olson.

#### References

1. E. J. Chaisson, "Gaseous Nebulae and Their Interstellar Environment," *Frontiers of Astrophysics*, edited by E. H. Avrett, Harvard University Press, 1976, p. 259-351.
2. W. J. Kaufmann, *Stars and Nebulas*, Freeman and Co., 1978.
3. J. B. Kaler, *Cosmic Clouds – Birth, Death, and Recycling in the Galaxy*, Scientific American Library, 1997.
4. Wen Zheng, C.R. O'Dell, "A Three-Dimensional Model of the Orion Nebula," *Astrophysical Journal*, Part 1, vol. 438, no. 2, January 1995, p. 784-793.
5. D. R. Nadeau, "Volume Scene Graphs," *submitted to IEEE Volume Visualization and Graphics Symposium 2000*.
6. B.A. Payne, A.W. Toga, "Distance Field Manipulation of Surface Models," *IEEE Computer Graphics and Applications*, vol. 12, no. 1, January 1992, p. 65-71.
7. S.F.F. Gibson, "Using Distance Maps for Accurate Surface Reconstruction in Sampled Volumes," *Proceedings of the IEEE Symposium on Volume Visualization*, 1998, p. 22-30.
8. K. Perlin, "An image synthesizer," *Proceedings of SIGGRAPH 85*, 1985. p.287-96.
9. M. Levoy, "Rendering of Surfaces from Volume Data," *IEEE Computer Graphics and Applications*, v. 8, no. 3, 1988, p. 28-37.
10. N. Max, "Optical Models for Direct Volume Rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 1, no. 2, June 1995, p. 99-107.
11. C.R. O'Dell, Kwan Wong, "Hubble Space Telescope Mapping of the Orion Nebula. I. A Survey of Stars and Compact Objects," *Astronomical Journal*, v.111, February 1996, p. 846-1002.
12. W. J. Kaufmann, *Universe*, 4<sup>th</sup> ed., Freeman and Co., 1994.
13. J. D. Genetti, "Multi-volume Rendering," *submitted to IEEE Volume Visualization and Graphics Symposium 2000*.