Implementation of Gooch et al's Non-Photorealistic Lighting Model For Automatic Technical Illustration

Kenton McHenry
University of Illinois Urbana-Champain
mchenry@uiuc.edu

May 6, 2003

1 Goals

The purpose of this project is to reproduce the results of [1] which describes a non-photorealistic lighting model for the purpose of rendering images simular to technical illustrations. Technical illustrations, unlike photos, are not concerned with realism but in conveying as much information as possible about the structure of objects. The phong lighting model commonly used is not suited for this purpose so an alternative model is presented.

2 The Phong Model

The phong lighting model is given by:

$$I = k_a + k_d \cos(N \cdot L) + k_s \cos(N \cdot L)^n$$

where I is a points intensity, k_d is its diffuse reflectence in RGB, k_a is the ambient illumination in RGB, k_s is the specular reflectance in RGB, n governs the extent of the specularity, N is the normal at the point (of unit length), and L is a unit vector pointing towards the light source. While this model is useful in aproximating real world lighting, it is not very effective in conveying all the information possible of the 3D structure. Figure-1 shows a human torso lit with the phong model.

Figure 1. Diffuse illumination only (ambience and specular set to zero)

Technical illustrations make particular use of silhouette edges and specular highlights to convey structural information. Figure-2 shows the above phong model with edges and highlights added. It is quickly seen that this is not satisfactory. The edges away from the light source (located up and to the right) are not visible because they blend in with the black of the object there. Highlights are also barley noticable due to the bright regions already present towards the light source.

Figure 2. Diffuse illumination (ambience set to zero), edges and highlights added

The edges and highlights do convey structural information as can be seen in Figure-3 which shows the same scene with only edges and highlights visible.

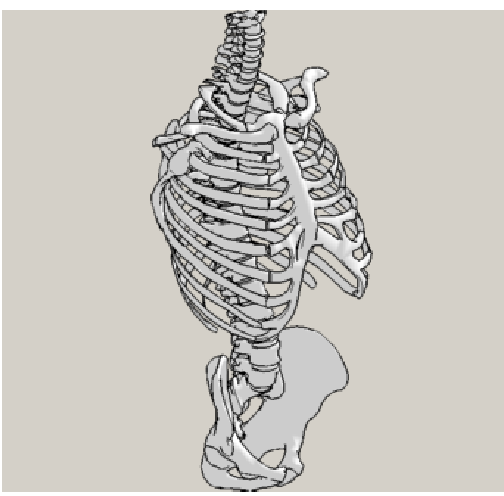

Figure 3. Edges and highlights only

We can attempt to overcome these problems within the phong model by limiting the diffuse component so that it can not reach white (making the highlights visible) and adding an ambient component so that it can not reach black (making the edges visible). An example of this is shown in Figure-4. The edges and highlights are visible but by limiting the range of intensities we have lost alot of the fine structure that was visible before and now blends into the ambient gray.

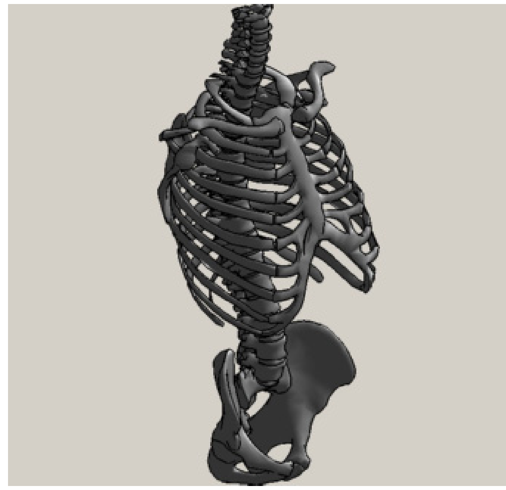

Figure 4. Diffuse illumination, with $0.5k_d$ and $k_a = (0.2, 0.2, 0.2)$, edges and highlights are more apparent

3 Tone-based Model

Technical illustrators make use of contrasting color temperatures to overcome the previous deficiency while still keeping black and white for edges and highlights respectivley. This leads us to our new model:

$$I = \left(\frac{1 + N \cdot L}{2}\right) k_{cool} + \left(1 - \frac{1 + N \cdot L}{2}\right) k_{warm}$$

where k_{cool} is a color of cool temperature such as blue, and k_{warm} is a color of warm temperature such as yellow. Rather than completley replacing the original diffuse color we use a linear combination:

$$k_{cool} = k_{blue} + \alpha k_d$$

$$k_{warm} = k_{yellow} + \beta k_d$$

where α and β are coeficients which should be given mid-level values within $[0, 1]$ in order to limit the range of the original diffused color preventing full black or white. k_{blue} is our cool color of the form $(0, 0, b)$, and k_{yellow} is our warm color of the form $(y, y, 0)$, where b and y are within the range $[0, 1]$ to limit the temeperature shift. Figure-5 shows the above object with the new lighting model. The finer detail is now visible aswell as the edges and highlights.

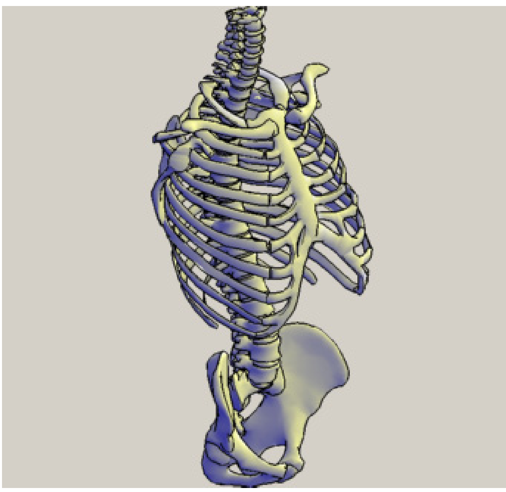

Figure 5. New lighting model with $\alpha = 0.2$, $\beta = 0.6$, $b = 0.4$, $y = 0.4$

4 Edges

We are concerned with three types of edges: silhouette, crease, and border (Figure-6). Creases are shared edges on faces where the angle is above some specifed threshold. Border edges are those which are on only one face. Both of these types can be found in a preprocessing step over the mesh and then stored. Silhouette edges on the other hand must be refound whenever the object is moved relative to the viewpoint.

A silhouette edge is defined as an edge that is shared by a front facing face and a back facing face. Finding all such edges involves testing every edge each time. A faster randomized alternative is described in [2]. Their process involves sorting the edges according to decreasing dihedral angle. Edges with a larger dihedral angle are more likley to be on the silhouette. So we are likley to get most of the long silhouettes by checking only a percentage of the edges with the largest dihedral angles. Once a silhouette is found we can walk along its neighboring edges to trace out the rest of the silhouette. It is often the case that silhouete edges in one viewpoint will also be on the silhouette in nearby viewpoints so we always recheck these edges. Even so this method tends to miss some of the silhouette edges on more complex objects.

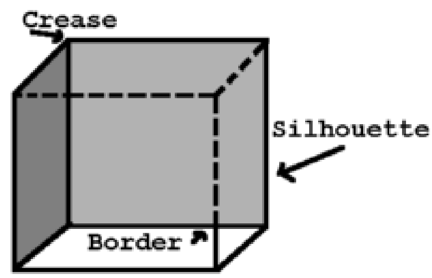

Figure 6. Different types of edges on a box with an open bottom

Once the edges are found we have decide on how to display them. We could simply render them all the same size and the same color such as Figure-7 but we could do better. More information can be conveyed to the viewer by drawing silhouette edges thicker (Figure-8). A better approach would be to draw lines in thicker that are closer to the viewpoint, such as Figure-9, giving the appearence of perspective. Finally creases can be drawn in white (Figure-10), giving the effect of a highlight on the object.

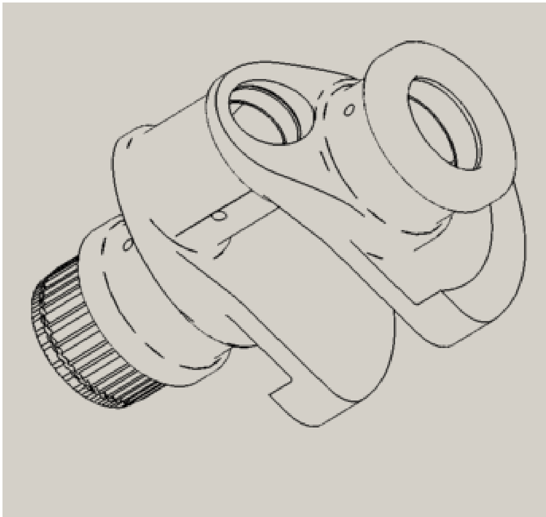

Figure 7. All lines have equal thickness

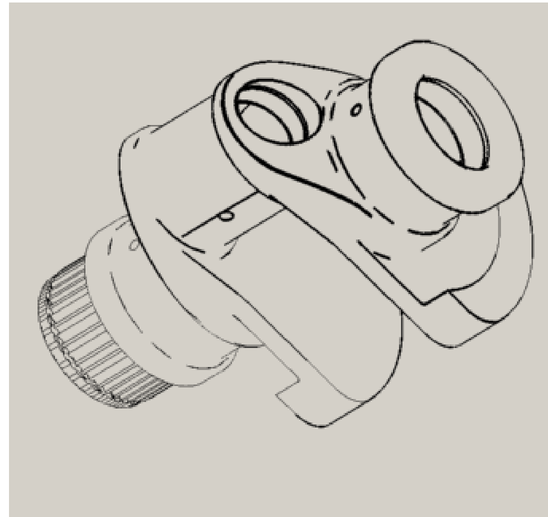

Figure 9. Lines closer to the viewpoint are thicker

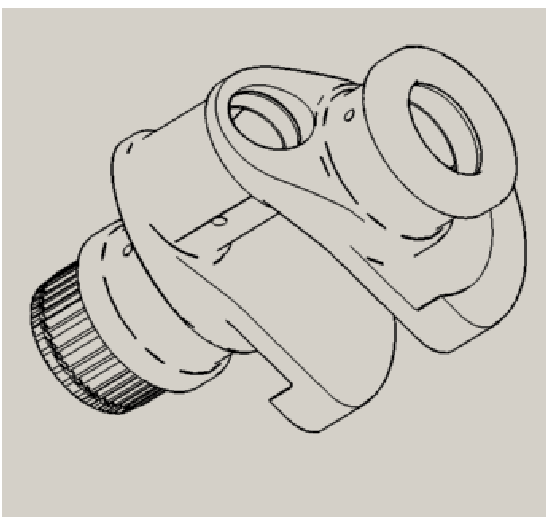

Figure 8. Silhouette lines are thicker

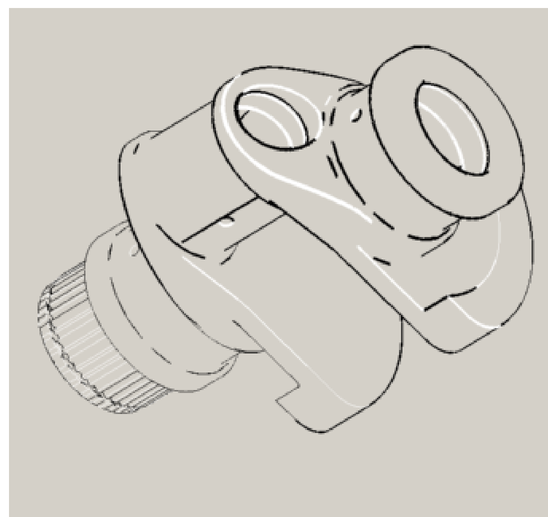

Figure 10. Creases are drawn white

5 Metal Shading

In technical illustrations metallic objects are drawn with very anistropic reflection (ussually only found on milled metal parts). This effect is created by placing bands of black and white lines along the axis of primary curvature. The black and white bands are constructed randomly using the formula:

$$I = b + (r * a)$$

where $b = -0.1$, r is a random number in the range $[0, 1]$, and $a = 0.4$. This formula is setup to favor black and white intensities. A 1D texture map with 20 values is then constructed. To make transitions between stripes somewhat smooth the texture is filtered with a 1-5-1 weighting scheme. Values are clamped to be between 0 and 1.

The direction of primary curvature at a vertex can easily be aproximated by examining a local neighborhood of vertices. We know the axis of primary curvature is in the direction opposite the normal at the vertex. We can find the distance along this direction with:

$$r = \frac{0.5}{\cos(\theta)}$$

where r is the radius off this axis, and θ is the angle between the unit ray of curvature and the inverse normal direction. We will use the angle between a ray formed by this axis and the vertex and some fixed ray to lookup the proper color in our 1D texture map. Finding a good reference ray is difficult without knowing all the vertices that share the same axis. So instead the light source direction is used. The above method for determining how far off the axis is not very reliable due to using only an aproximate value of primary curvature. However this does not matter when using a light source as our fixed reference ray when it is at infinity. In this case the magnitude off the normal does not effect the angle since all the rays are the same.

Figure-11 shows a metal shaded object. The idea that this object is made of metal is definatley conveyed by the stripes. Figure-12 shows the same object with the metal shading and the previous tone based shading. Though not as convincing as the previous image, it contains the important structural information provided by the tones while still giving the impression of being made of metal.

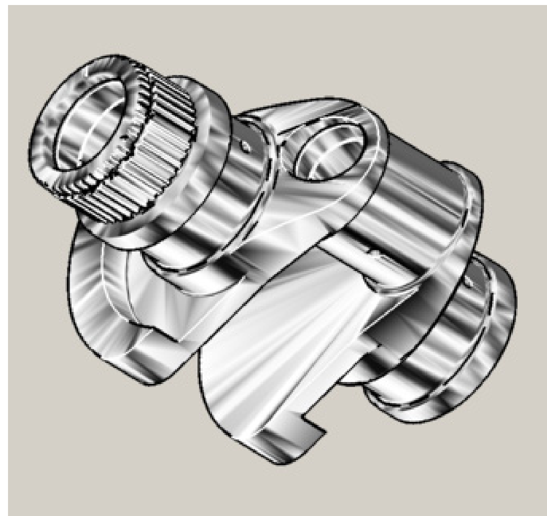

Figure 11. Metal shaded object

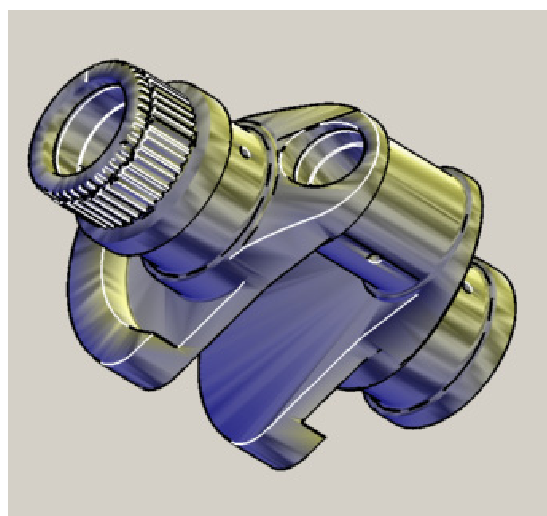

Figure 12. Metal and tone shaded object

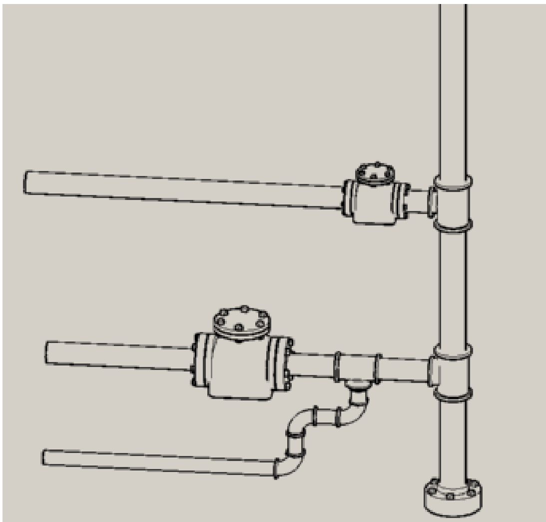

Figure 13. Edges only

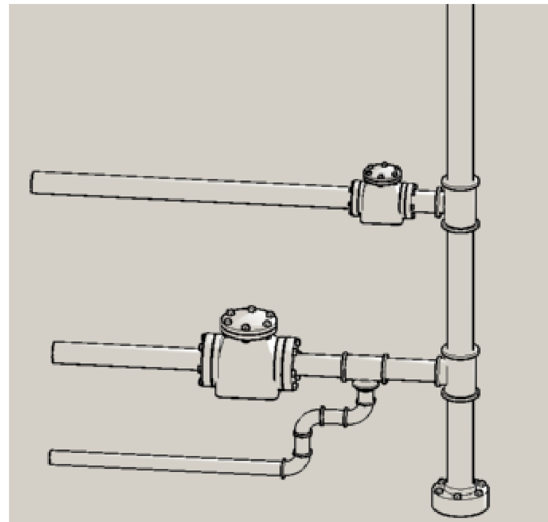

Figure 15. Edges and highlights

Figure 14. Highlights only

Figure 16. Diffuse shading

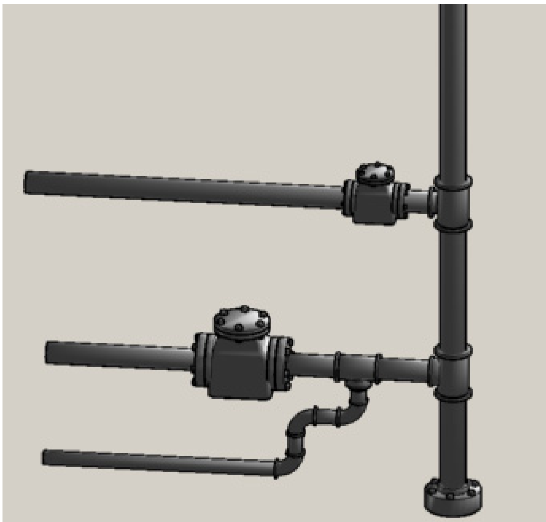

Figure 17. Limited diffuse shading with ambient illumination to make edges and highlights visible

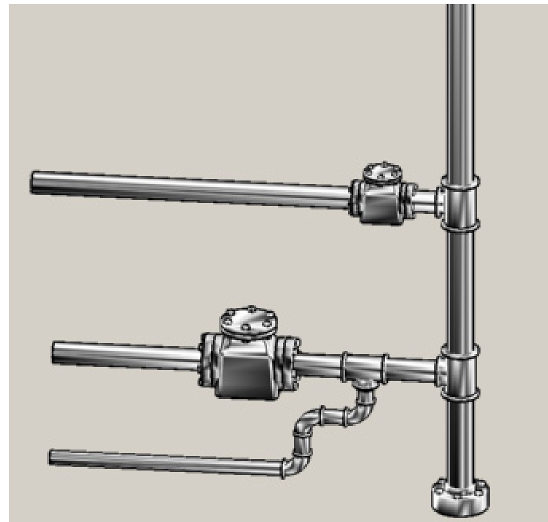

Figure 19. Metal shaded

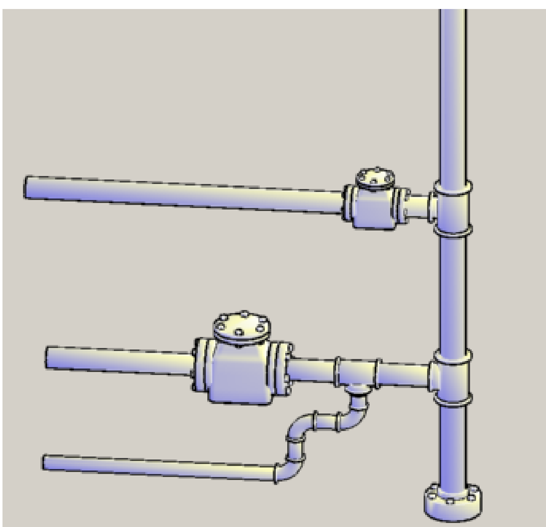

Figure 18. Tone shaded

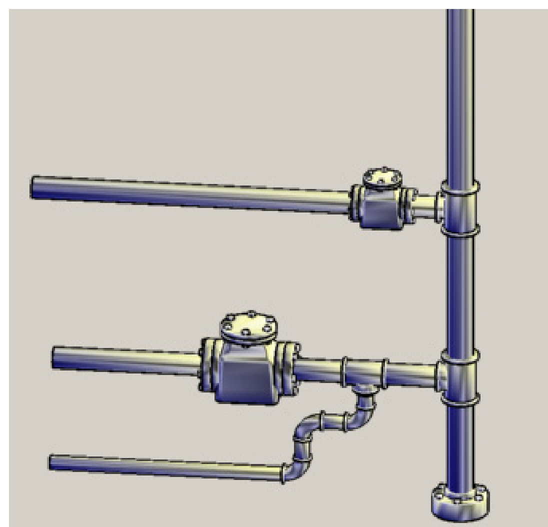

Figure 20. Metal and tone shaded

6 Remarks

The experiments show that using color temperature is useful in conveying structural information. As can be seen in Figure-17 the base pipe is a solid gray, giving us no structural information. In Figure-18 we see the color variation indicating the change in normal direction. The randomized silhouette algorithm performed poorly with complex objects, leaving gaps in the silhouette, and wasn't that much of a speed up. Finaly the metal shading technique effectivley conveyed the impresion of a milled metal surface.

References

[1] Gooch, Gooch, Shirley, Cohen, *A Non-Photorealistic Lighting Model For Automatic Technical Illustration*, SIGGRAPH 1998.

[2] Markosian, Kowalski, Trychin, *Parametric Equation of a Sphere and Texture Mapping*, SIGGRAPH 1997.