Whimsical Night: Artistic Control of a Physically Based Rendering Model

Eric Chen

Figure 1: *Whimsical Night.* 1920 × 1080p, 4096 spp, 9 hours

1 Introduction

For my final project, I wanted to explore the dichotomy of modeling physical phenomena and creative abstraction. I decided to render Eevee, a Pokemon, with realistic fur and lighting, but in an imaginary world. To create a physically based renderer with artistic control, I first implemented several features to add realism: texture mapping, environment mapping with hierarchial sampling, and a thin lens camera. Finally, I spent most of the time implementing the hair model proposed in Chiang et al. (2016), thanks to a great supplemental chapter in the PBR book Pharr et al. (2016). What makes this paper different from prior works is that it introduces an intuitive way to control hair roughness and reflectance to physical parameters. The method has been adopted by Disney artists and animators in production.

2 Texture Mapping

Texture mapping is a cheap way to bring life into a scene. I store the textures in mipmaps, and use them as a lookup table. I show the effect of texture mapping in Figure 2 and Figure 3.

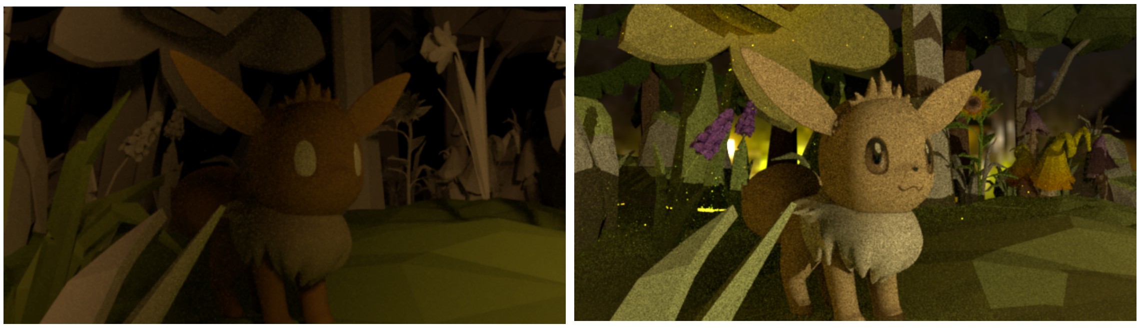

Nori Starter — Nori with textures and environmental lighting

Figure 2: Adding texture mapping and environment mapping show a large improvement over the Nori base code.

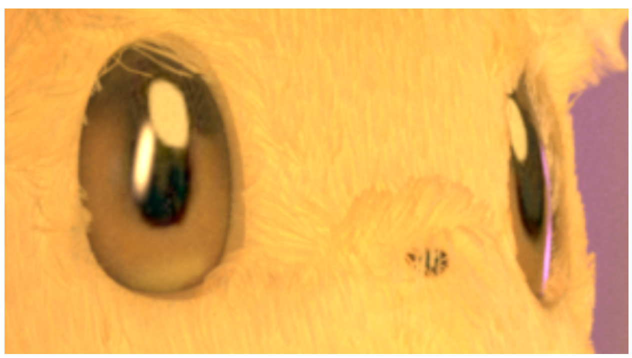

Figure 3: Eevee's eyes show a combination of fake reflection from the texture map, and real reflection from the microfacet material.

3 Environment Mapping and Hierarchial Sampling

I implemented a hierarchial sampling strategy for the environment map. Hierarchial sampling is a fast way to importance sample with respect to luminance. Thankfully, I did not have many issues with high variance fireflies like what other groups that did not do importance sampling reported. Please see Section 6 for ablations on sampling rate. I show the effect of two different environment maps in Figure 4.

4 Thin Lens Camera Model

The thin lens model causes defocus blur by averaging sampled points within a radius of the rendered ray. This adds realism to the scene since it simulations the effect of taking a picture with a real camera. I show the effect of different aperture radii in Figure 5.

5 Rendering Hair

I implemented the hair model proposed in Chiang et al. (2016). The hair was sculpted in Blender, and exported as cylinder triangle meshes to Nori. Unfortunately, this does limit the realism of the hair, as the scattering equations are designed to work on cylinders.

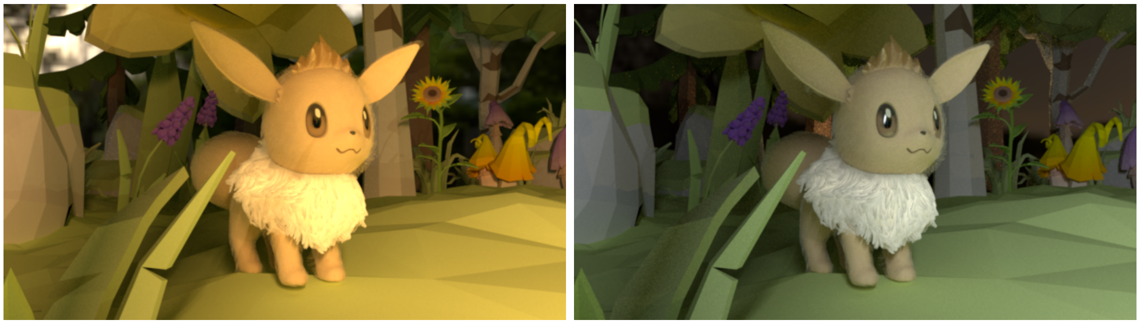

Day Night

Figure 4: The Eevee scene under two different environment maps

0.00 (Pinhole) 0.02 0.02

Figure 5: The effect of the thin lens at different aperture radii.

I outline the hair model below. Rendering hair is a longstanding problem in computer graphics. Kajiya & Kay (1989) made one of the first attempts to render hair, and the model was improved upon significantly by Marschner et al. (2003). What made Marschner et al. (2003) unique at the time is that the authors took a very scientific approach to rendering hair by measuring the scattering of individual hair fibers, validating their model.

5.1 Hair Anatomy

Hair consists of three major parts: the cuticle, which is the outer layer, the cortex, which consists of 90% of hair volume, and the medulla, which is a thick core. The cuticle consists of scales angled at about 2°, and can be modeled as a rough dielectric. The cortex absorbs most light, and the medulla causes scattering. In animals, the medulla is particularly thick. Naturally, this anatomy corresponds to three scattering phenomena: longitudinal scattering, which is scattering along the length of the hair, azimuthal scattering, which is scattering on the circumference, and absorbtion.

5.2 The Three Lobes

Marschner et al. (2003) represent scattering through hair as three simple interactions: R (reflection), TT (transmission, transmission), and TRT (transmission, reflection, transmission). While this model is simple, it results in high energy loss, especially when hair has low absorbtion, as in lighly colored hair. Nonetheless, this model integrates nicely into a path tracer.

5.3 The Hair Model

In this section I describe the equations for scattering.

Longitudinal Scattering. The model for longitudinal scattering used in Chiang et al. (2016) is from d'Eon et al. (2011). Given incoming and outgoing angles θ_o, θ_i, the scattering function is

$$M_p(\theta_o, \theta_i) = \frac{1}{2v \sinh(1/v)} e^{\frac{-\sin\theta_i \sin\theta_o}{v}} I_0\left(\frac{\cos\theta_o \cos\theta_i}{v}\right), \tag{1}$$

where v is a roughness parameter and I_0 is a modified Bessel function of the first kind.

Azimuthal Scattering. Prior work such as d'Eon et al. (2011) use Gaussian quadrature to model azimuthal scattering. Chiang et al. (2016) instead use a logistic distribution, which is simpler. Given an angle ϕ and roughness parameter s, the distribution over $[-\pi, \pi]$ is

$$D(\phi) = \frac{1}{\frac{1}{1+e^{-\pi/s}} - \frac{1}{1+e^{\pi/s}}} \frac{e^{-\phi/s}}{s(1+e^{-\phi/s})^2}, \tag{2}$$

Absorption. Absorption in hair is modeled as done in classical volume rendering, with a few caveats. However, I will omit this for brevity. For a detailed explanation please see Pharr et al. (2016).

5.4 Artistic Controllability

A major contribution of Chiang et al. (2016) is that they map linear roughness parameters β_M (for longitudinal scattering), and β_N (for azimuthal scattering), to seemingly cryptic variables

$$v = (0.726\beta_M + 0.812\beta_M^2 + 3.7\beta_M^{20})^2 \tag{3}$$

$$s = 0.265\beta_N + 1.194\beta_N^2 + 5.372\beta_N^{22}. \tag{4}$$

Notice how large the powers of the polynomials are. The authors state this corresponds to artistic perception, and works well in practice.

Furthermore, the authors present a way to map perceived albedo C to the absorption parameter:

$$\sigma_a = (\ln C/(5.969 - 0.215\beta_N + 3.532\beta_N^2 - 10.73\beta_N^3 + 5.574\beta_N^4 + 0.245\beta_N^5))^2. \tag{5}$$

I show a comparison between the diffuse BSDF and the hair BSDF on our scene in Figure 6. Notice how energy conservation is much better with the hair BSDF.

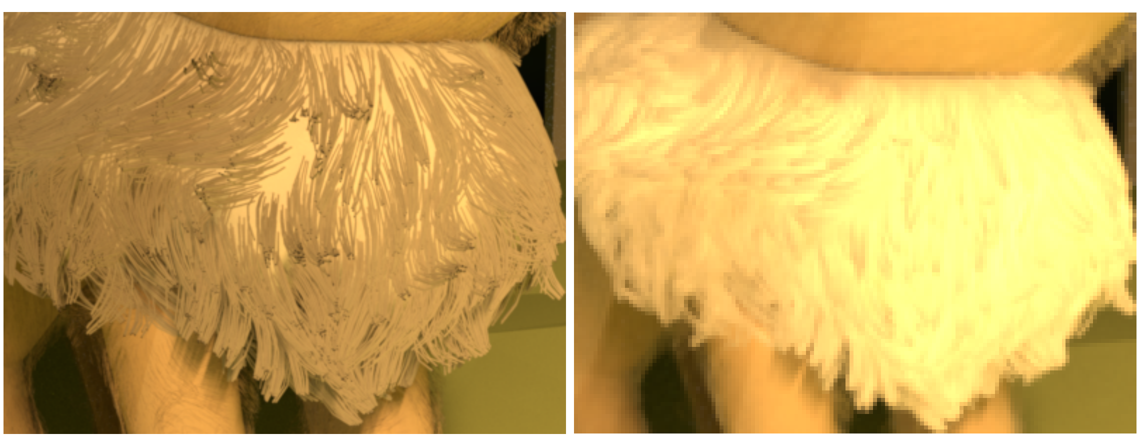

Diffuse BSDF | Hair BSDF

Figure 6: Using the diffuse BSDF, the hair looks sparser, due to the reflection of individual fibers. With the hair BSDF, it acts more like a volume. The hair BSDF shows much higher reflectance. Energy is also better conserved. Sorry for the difference in resolution.

6 Implementation Details

Final renders were done on a 32 core CPU. The submission photograph had 66 million triangles, most of which are from the hair. I show the effect of sampling rate in Figure 7.

1 spp, 8 sec | 16 spp, 2 min

64 spp, 8 min | 128 spp, 14 min

512 spp, 1 hour | 4096 spp, 9 hours

Figure 7: I show the relationship between sampling rate and noise. After 128 spp most noise is not noticable.

7 Conclusion

I tried to model a variety of artistic effects using physically based rendering. Overall, the project was a lot of fun. I learned a lot about both physics and rendering.

Grading Notes by Steve:

- Centered on fur but has several other components
- Added textures
- Env map sampling using hierarchical warping
- Thin lens camera model

- Hair is the most extensive part
- Discussion of hair anatomy and optics
- Implementation is somehow applied to triangle mesh fibers
- gets a nice volumetric look but unclear whether benefits of hair BSDF are there
- Overall a very nice and finished looking image

Report

Echoes the thorough and well-prepared presentation. You spend a lot of time talking about the Chiang et al. hair model, but don't show any results from your implementation that would indicate whether or not the implementation works. The hair does end up looking soft, which is good, though I'm quite confused by the comparison showing diffuse with an odd glow behind the fibers and hair looking quite emissive (possible normalization issue?), but also maybe with an overall blurrier image.

The contest image is impressive and nicely refined/tuned; it shows off the various new features effectively.

4.0 / 4.0

References

Matt Jen-Yuan Chiang, Benedikt Bitterli, Chuck Tappan, and Brent Burley. A Practical and Controllable Hair and Fur Model for Production Path Tracing. *Computer Graphics Forum*, 2016. ISSN 1467-8659. doi: 10.1111/cgf.12830.

Eugene d'Eon, Guillaume Francois, Martin Hill, Joe Letteri, and Jean-Marie Aubry. An energy-conserving hair reflectance model. In *Proceedings of the Twenty-Second Eurographics Conference on Rendering*, EGSR '11, pp. 1181–1187, New York, NY, USA, 2011. Association for Computing Machinery. doi: 10.1111/j.1467-8659.2011.01976.x. URL https://doi.org/10.1111/j.1467-8659.2011.01976.x.

J. T. Kajiya and T. L. Kay. Rendering fur with three dimensional textures. In *Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '89, pp. 271–280, New York, NY, USA, 1989. Association for Computing Machinery. ISBN 0897913124. doi: 10.1145/74333.74361. URL https://doi.org/10.1145/74333.74361.

Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. *ACM Trans. Graph.*, 22(3):780–791, jul 2003. ISSN 0730-0301. doi: 10.1145/882262.882345. URL https://doi.org/10.1145/882262.882345.

Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation (3rd ed.)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, November 2016. ISBN 9780128006450.