Predictable and Targeted Softening of the Shadow Terminator

Priyamvad Deshmukh
DreamWorks Animation
priyamvad.deshmukh@dreamworks.com

Brian Green
DreamWorks Animation
brian.green@dreamworks.com

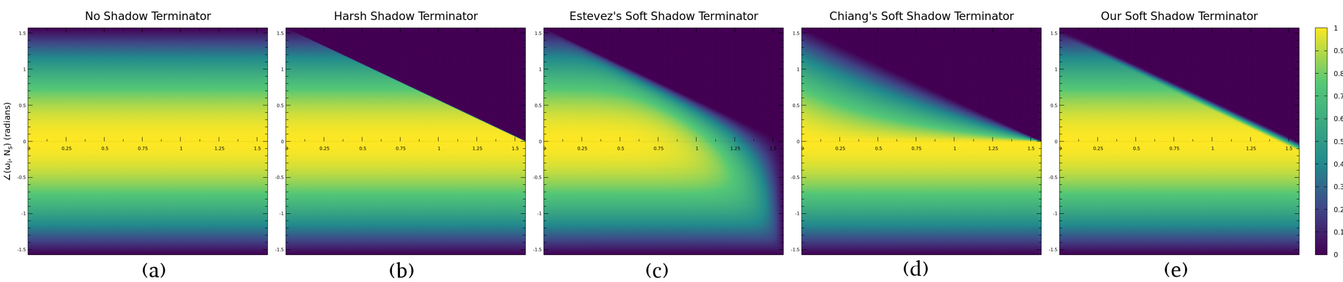

Figure 1: (a) Smooth cosine falloff with no shadow terminator, (b) harsh-edge artifact from discarding light beyond the geometric hemisphere (c) the soft terminator from [Estevez et al. 2019] (d) the soft terminator from [Chiang et al. 2019]. (e) our soft shadow terminator.

ABSTRACT

A well known artifact in production rendering from the use of shading normals is the shadow terminator problem: the abrupt interruption of the light's smooth cosine falloff at geometric horizons. Recent publications introduced several ad-hoc techniques, based loosely on microfacet theory to deal with these issues. We show that these techniques can themselves introduce artifacts and suggest a new technique that is an improvement in many situations. More importantly we introduce a framework for analyzing these different techniques so artists and researchers can choose appropriate solutions and more reliably predict and understand expected results.

CCS CONCEPTS

• **Computing methodologies → Rendering**; **Ray tracing**;

KEYWORDS

ray tracing, shadow terminators

ACM Reference Format:
Priyamvad Deshmukh and Brian Green. 2020. Predictable and Targeted Softening of the Shadow Terminator. In *Proceedings of SIGGRAPH*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3388767.3407371

1 INTRODUCTION

Non-geometric shading normals are a common practice in production path-tracing. Texturing artists employ this tool to create complex structures on relatively simple geometry. While this improves the look with added detail, it also leads to artifacts caused by physically impossible situations. To avoid light leaks, directions outside the geometric hemisphere are discarded. But this leads to abrupt black fringe arifacts, solutions for which are explored in [Chiang et al. 2019] and [Estevez et al. 2019]. We evaluate these solutions based on a common framework we propose below.

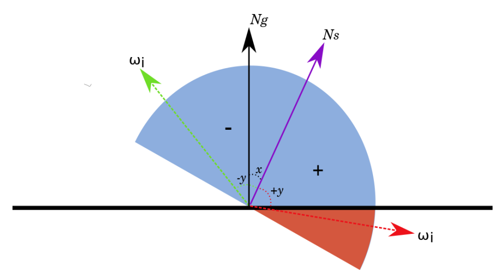

Figure 2: Terminating the smooth cosine hemisphere to avoid light leaks results in a harsh shadow terminator.

2 FRAMEWORK

Let us establish the notation we will use throughout the paper to describe the shadow terminator. The angle between the shading normal (Ns) and the geometric normal (Ng) can deviate between $[0, \frac{\pi}{2}]$ and is denoted by x. This forms the horizontal axis in Figure 1. Within the shading hemisphere, the light direction (ω_i) can range from $[\frac{-\pi}{2}, \frac{\pi}{2}]$ with respect to Ns. The positive or negative direction is chosen based on the geometric normal. The directions ω_i towards Ng are labeled negative and the ones away from Ng are labeled positive (Figure 2). This is the vertical axis in Figure 1. This angle between Ns and ω_i is denoted as y. With this arrangement, the angle between Ng and ω_i comes out to be $(x + y)$. To avoid light leaks, we have to cull all incoming light direction when $(x+y) \geq \frac{\pi}{2}$

creating a stark line that depicts the harsh shadow terminator in Figure 1(b).

With these notations in place, we can show that the existing techniques for softening the shadow terminator can all be formulated in the same framework and plotted accordingly. For example, the [Chiang et al. 2019] shadow terminator is equivalent to the following and depicted in Figure 1(d):

$$G = \min\left[1, \frac{\cos(x+y)}{\cos(x)\cos(y)}\right] \tag{1}$$

Similarly, the terminator from [Estevez et al. 2019] can be formulated and is depicted in Figure 1(c):

$$G = \frac{2}{1 + \sqrt{1 + (\frac{1}{8}\tan^2(x)\tan^2(x+y)}} \tag{2}$$

While loosely based on inspiration from microfacet theory, both these factors soften the shadow terminators in a non-intuitive manner, as can be seen in Figure 1(c) and 1(d). More importantly, they also have an influence away from the shadow terminator, dampening highlights from the BRDF. The visual effect can be more or less severe based on the BRDF and the outgoing vector. We can verify that, as claimed, the technique in [Chiang et al. 2019] performs better than the one in [Estevez et al. 2019] in the valid geometric hemisphere when $y \leq 0$, by not dampening the highlights at all. However, looking at Figure 1(d), what is also clear is that it excessively dampens highlights in the visible shading hemisphere when $y > 0$ and $(x+y) \leq \frac{\pi}{2}$. Moreover, small deviations from Ng are penalized heavily in [Chiang et al. 2019]. This creates unexpected loss of detail as seen in Figure 3 (b) and (c).

3 OUR APPROACH

While designing our approach to soften the shadow terminator, our priorities were simple, (1) be predictable and (2) do not affect any other highlights. To this end, we analyzed the plot which forms a straight line $(x+y) \geq \frac{\pi}{2}$ in Figure 1(b). We then used a smooth step function as the shadowing term on top of the regular cosine factor in the rendering integral. We designate the smoothing region as $[\cos(\frac{\pi}{2}), \cos(\frac{\pi}{2} - \alpha)]$, or $[0, \sin(\alpha)]$, where α is the angle at which we do not want the shadow terminator to have any effect. This results in our terminator:

$$t = \min\left[\max\left[\frac{\cos(x+y)}{\sin(\alpha)}, 0\right], 1\right] \tag{3}$$

Similar to [Chiang et al. 2019], we use a cubic Hermite interpolator to avoid any sharp discontinuities:

$$G = 3t^2 - 2t^3 \tag{4}$$

In practice, we want a smoother falloff when the deviation from geometric normal is large. So we blend between two thresholds as is evident in the code snippet we provide below.

```
float shadowTerminator(const Vec3f& Ng, const Vec3f& Ns, const Vec3f& wi)
{
    const float alpha = 0.05;
    float d = lerp(sin(alpha+0.1), sin(alpha), dot(Ns, Ng));
    float t = max(0.0f, min(1.0, dot(Ng, wi)/d));
    float G = t*t*(3.0 - 2.0*t);
    return G;
}
```

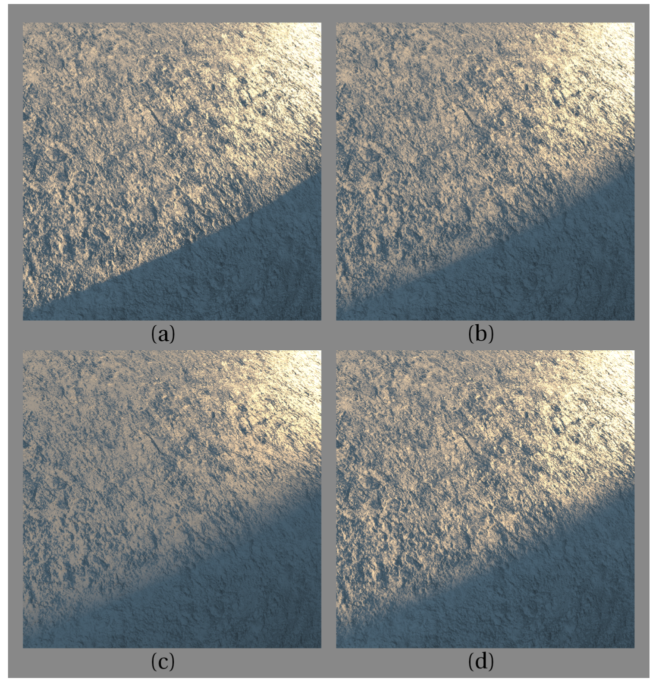

Figure 3: Various shadow terminators: (a) the artifact, (b) the soft terminator from [Estevez et al. 2019] (c) the soft terminator from [Chiang et al. 2019] (d) our soft shadow terminator.

This terminator can be visualized in Figure 1(e) and the renders in Figure 3(d). As is clear, our soft shadow terminator only affects the harsh-shadow edge and preserves all the remaining desired highlights from shading normals.

4 CONCLUSIONS AND FUTURE WORK

We have introduced a framework for analyzing techniques to soften the shadow terminator. We also provide a new, more predictable technique that preserves valid BRDF highlights. Since many plausible solutions exist for this problem, we hope that our framework and notation for comparing shadow terminators can provide a common language when evaluating future development. In the future, we would like to explore solutions to compensate for energy loss that results from the use of shading normals.

ACKNOWLEDGMENTS

This work was implemented by the MoonRay development team at DreamWorks Animation. Rigorous testing, feedback, and patience from lighting artists like Pablo Valle and Iveth Bueno were key to success.

REFERENCES

Matt Jen-Yuan Chiang, Yining Karl Li, and Brent Burley. 2019. Taming the Shadow Terminator. *ACM SIGGRAPH 2019 Talks* 1, Article 71 (July 2019). DOI: http://dx.doi.org/10.1145/3306307.3328172

Alejandro Conty Estevez, Pascal Lecocq, and Clifford Stein. 2019. *A Microfacet-Based Shadowing Function to Solve the Bump Terminator Problem.* Apress, Berkeley, CA, 149–158. DOI: http://dx.doi.org/10.1007/978-1-4842-4427-2_12