Specular Lobe Aware Upsampling Based on Spherical Gaussians

Yusuke Tokuyoshi*
Square Enix Co., Ltd.

Introduction This poster introduces a novel weighting function of bilateral upsampling for specular surfaces. High-resolution rendering with expensive shaders (e.g., global illumination) is a considerable problem for real-time applications such as video games. Therefore, some bilateral upsampling based methods were proposed to alleviate the computational burden. Spatio-temporal upsampling [Herzog et al. 2010] is also proposed to generate more sample pixels than spatial bilateral upsampling by reusing past frames. The main challenge of these upsampling techniques is to optimize a weighting function in order to estimate a pixel value by appropriately prioritizing samples. The weighting function commonly evaluates a similarity of pixel values. For example, a surface normal continuity is generally used for evaluating a geometric similarity, which is well suited for diffuse surfaces because the shading results depend on the normal vectors. However, for specular surfaces, we have to take into account not only the normal vectors but also the eye directions and the specular sharpness to generate accurate results, since specular lobes are determined by them.

We propose an efficient weighting function based on a specular lobe similarity. This function is simple and has no additional storage cost. Moreover, it has no user-specified parameters. Thus it can be easily integrated with general upsampling techniques.

Our Weighting Function A weighting function should evaluate a similarity of the reflected radiance between the current pixel i and the sampled pixel j, which is the inner product of the incident radiance and the reflection lobe. Assuming that the incident radiance distributions of two pixels are identical, the difference of the reflected radiance is determined by only the reflection lobes. Therefore, we introduce a weighting function $w_{i,j}$ representing the specular lobe similarity as given by computing the inner product of the two specular lobes. However, the shape of the specular lobe depends on the BRDF model. Furthermore, there is not always an analytical solution of the inner product.

In this poster, we approximate the lobes by using spherical Gaussians (SGs). As described in [Wang et al. 2009], a specular lobe can be approximated with an SG $G(\boldsymbol{\omega})$ where $\boldsymbol{\omega}$ is the incident direction of light. The inner product of two SGs is analytically obtained. Hence, we define the weighting function as the inner product of the normalized SGs:

$$w_{i,j} = \frac{G_i(\boldsymbol{\omega})}{\|G_i(\boldsymbol{\omega})\|} \cdot \frac{G_j(\boldsymbol{\omega})}{\|G_j(\boldsymbol{\omega})\|}. \quad (1)$$

The range of this weighting function is [0, 1]. When the two lobes are same, $w_{i,j} = 1$.

Applications and Conclusion We here employ voxel cone tracing [Crassin et al. 2011] with spatio-temporal upsampling as an experimental example since voxel cone tracing is known to be computationally expensive to render specular surfaces.

Figure 1 shows the comparison of a general normal based weighting function and the proposed weighting function. In the left images using the normal based weighting function, there are some blurring and flickering artifacts due to estimation errors. On the other hand, our method enables higher-quality upsampling without parameter tuning. Our approach adaptively reduces blurring and flickering artifacts depending on specular sharpness and eye directions, in contrast to the normal based weighting function does not.

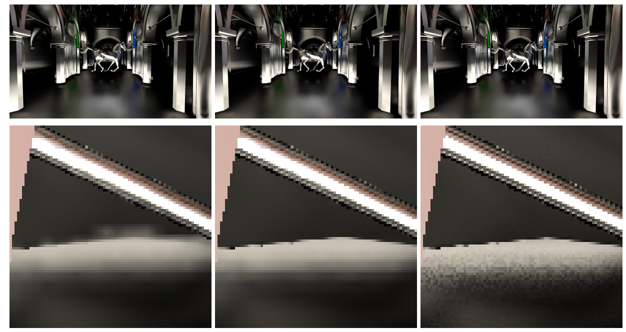

normal based	proposed	naïve rendering
2.3 ms	2.5 ms	11 ms

Figure 1: *Upper: the specular indirect illumination results via voxel cone tracing. Lower: close-ups of the upper images (brightness×4). The left and middle images are upsampled from* 480×270 *to* 1920×1080 *pixels. The right images are rendered with naïve per-pixel cone tracing. The high-frequency noise is due to stochastic sampling for ray marching. The BRDF is the Blinn-Phong model (phong exponent: 1023). The normal based weighting function has a Gaussian distribution (the variance parameter: 0.005) for this experiment. The computation time of cone tracing and upsampling is only 2.3 ms (left) and 2.5 ms (middle), while naïve per-pixel cone tracing is 11 ms (right) (GPU: AMD Radeon HD 6990). In addition, our method (middle) produces closer images to the truth.*

Our approach is limited to isotropic BRDFs which can be approximated by a single SG. However, this is not a problem for voxel cone tracing which also approximates BRDFs by a Gaussian lobe. In addition, we can introduce anisotropic BRDFs using several SGs with sacrificing time.

Our weighting function is also applicable for geometry aware blurring (i.e., specular lobe aware blurring). In future work, we would like to investigate the effectiveness of such bilateral filtering.

References

CRASSIN, C., NEYRET, F., SAINZ, M., GREEN, S., AND EISEMANN, E. 2011. Interactive indirect illumination using voxel cone tracing. *Comput. Graph. Forum 30*, 7, 1921–1930.

HERZOG, R., EISEMANN, E., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2010. Spatio-temporal upsampling on the gpu. In *Proc. of I3D 2010*, 91–98.

WANG, J., REN, P., GONG, M., SNYDER, J., AND GUO, B. 2009. All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Trans. Graph. 28*, 5, 133:1–133:10.

*e-mail:tokuyosh@square-enix.com

SIGGRAPH 2013, July 21 – 25, 2013, Anaheim, California.

ACM 978-1-4503-2261-4/13/07