(Talk 0263) Smoother Subsurface Scattering Supplemental Notes

Ivan Neulander
Rhythm & Hues Studios

April 23, 2009

These supplemental notes provide additional sample renderings and timing information.

Our teapot is 1.2 units tall. It was rendered with a mean scatter distance ℓ_u of 0.2 units for red, 0.1 for green, 0.05 for blue, except where noted otherwise. For smoothing we used the function $w(r) = \max(1 - r/\alpha, 0)$ for local weighting, where the smoothing radius α is four times the radius of a circle whose area equals the average surface area per irradiance sample point. Without smoothing, objectionable levels of noise are evident in the green and especially blue components of all but the ground truth frames.

All images were rendered on an AMD Athlon64 X2 2.8 GHz, 4GB RAM. All still images were rendered at 1024x1024 pixel resolution with one shading sample per fragment. The attached video was rendered at 640x480 resolution with one shading sample per fragment, single-thread execution.

For Figure 1 in the Abstract, the CPU times were as follows.

- No smoothing (23k irradiance cache points): 413s
- Langlands & Mertens smoothing (23k points): 422s
- Our smoothing (23k points): 411s
- Ground truth (1 million points, no smoothing): 657s

Figure 3 in the Abstract took:

- No smoothing (10k points): 407s
- Our smoothing with integral interpolation (10k points): 427s
- Ground truth (1 million points, no smoothing): 1339s

The attached video compares the smoothing methods for animated geometry at fixed ℓ_u and sample density settings. Average CPU times per frame were:

- No smoothing (23k irradiance cache points): 220s
- Langlands & Mertens smoothing (23k points): 224s
- Our smoothing (23k points): 228s
- Control (230k points, no smoothing): 332s

These statistics confirm that our smoothing implementation adds minimal overhead and saves time relative to raising the irradiance sample density. While the savings relative to the brute-force ground truth case are modest, this is typical for ℓ_u values high enough to yield soft shadow edges. As the mean free path diminishes, smoothing produces greater savings, and our solution converges with that of Langlands and Mertens, since all shadow edges become increasingly hard. But at higher ℓ_u values, our method preserves the accuracy of softer shadow edges, allowing it to be used safely over a wide range of mean free paths.

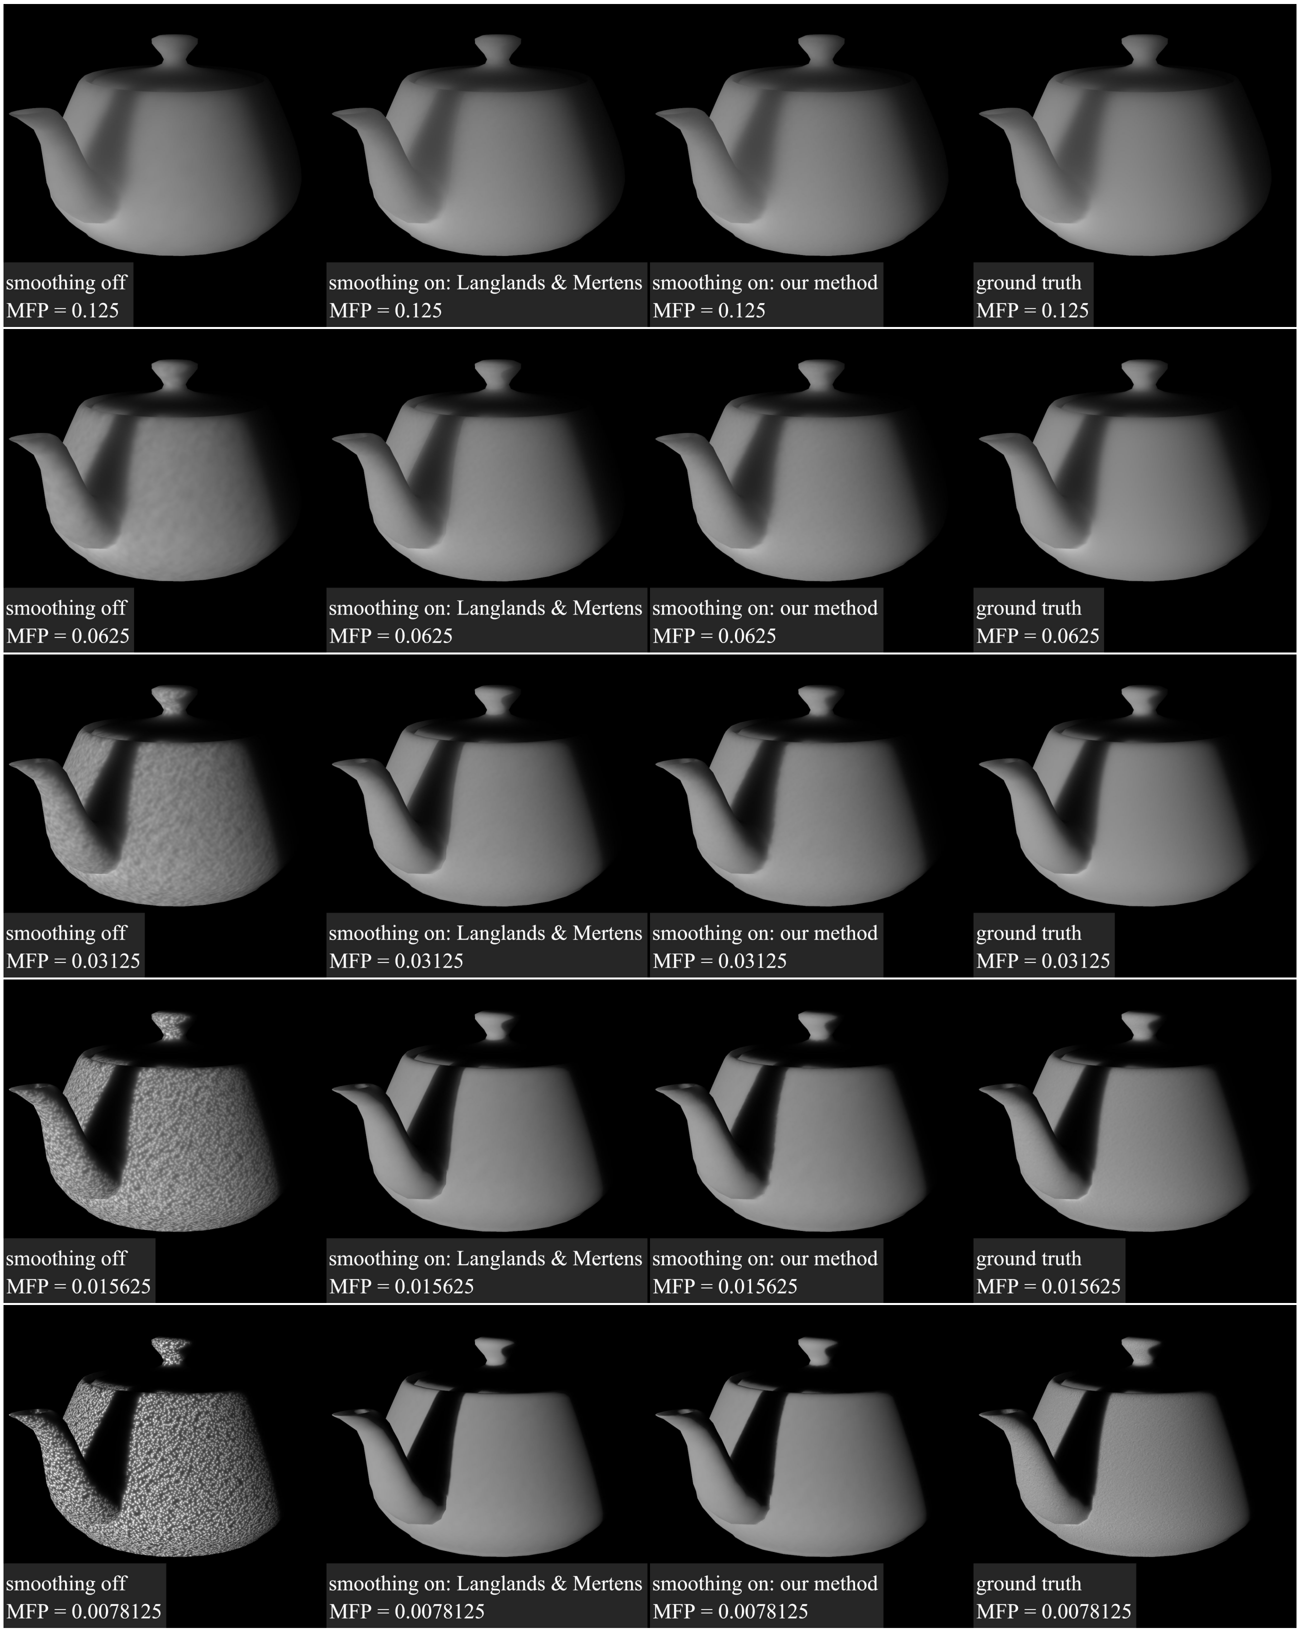

Figure 1: These renderings compare the smoothing methods over varying monochrome ℓ_u, at fixed sample densities. Columns from left to right: no smoothing (80k irradiance points, 235s per frame), Langlands & Mertens smoothing (80k points, 233s), our smoothing (80k points, 230s), ground truth (one million points with no smoothing, 335s). Note minor noise in “ground truth” for bottom two rows: one million cached points is evidently insufficient for the lowest mean free paths.

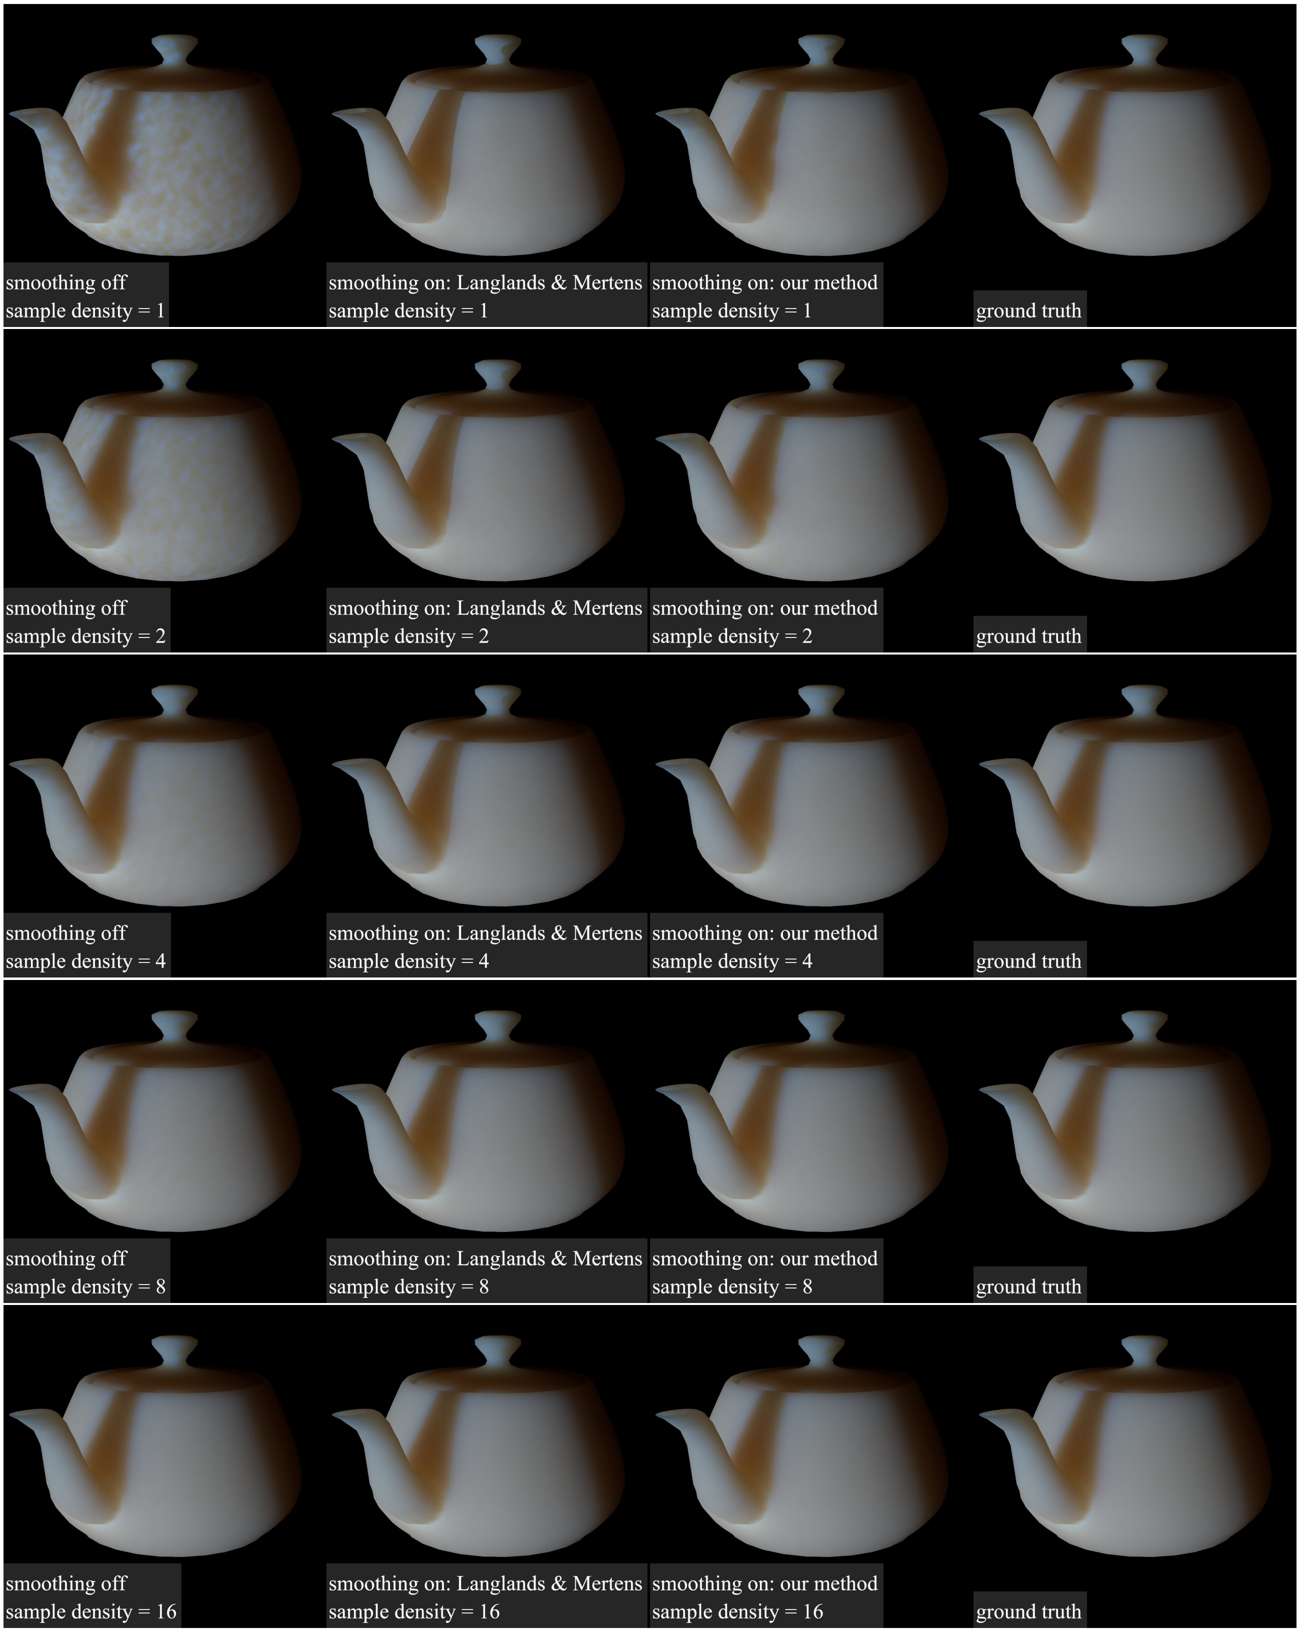

Figure 2: These renderings compare smoothing methods over variable sample densities (ranging from 2k to 1 million irradiance points) at a fixed RGB ℓ_u of (0.2, 0.1, 0.05). Columns from left to right: no smoothing, Langlands & Mertens smoothing, our smoothing, ground truth (identical in all rows). Average frame times were consistent with the figures above.