Bayesian Collaborative Denoising for Monte Carlo rendering
Additional materials

Supplementary results: compared methods

For each scene, we compare:

- ARNLMF (Adaptive rendering with non-local means filtering) [Rousselle2012]
- RDFCI (Robust Denoising using Feature and Color Information) [Rousselle2013]
- WLR (Weighted Local Regression) [Moon2014]
- LBF (Learning Based Filtering) [Kalantari2015]
- RHF (Ray Histogram Fusion) [Delbracio2014]
- ours (Bayesian Collaborative Denoising)

Note that ARNLMF and WLR are using adaptive sampling while ours performs a pure denoising and does not require triggering additional samples from the renderer.

Supplementary results: timings

All computations are performed using a 4-core Intel Xeon processor at 2.80 GHz, and are always taking advantage of these 4 cores through CPU-parallelization. ARNLMF, RDFCI, WLR and LBF additionally use CUDA for GPU-parallelization (for filtering, not for rendering).
All implementations use PBRT as rendering engine. ARNLMF, RDFCI, WLR and LBF are coded as PBRT extensions, and perform rendering, filtering and sometimes adaptive sampling in the same executable. RHF and ours use simple customizations of PBRT to produce additional inputs, then perform the filtering in a separate program. Hence, for timings of RHF and ours, we measure the filtering time and add it to the one that was needed to produce the input "Simple PBRT rendering".

Supplementary results: denoising measure

For each method, we measure the relative Mean Square Error (relMSE) and Structural Similarity (ssim) between the result and the groundtruth (generated with a much higher number of samples per pixel). A good denoising has a low relMSE (minimum 0) and a high ssim (maximum 1). All comparisons are equal-sample. Even the methods that perform adaptive sampling target the same average number of samples per pixel (spp). Note that some of the methods, like LBF, are able to provide nice results with a very low number of spp. But here, we assume a more production-like context where we can afford a reasonably large amount of samples per pixel and seek very high quality results.

Supplementary results: scenes

- SanMiguelDof (256spp) Relatively complex scene with depth of field
- SponzaDir (256spp) Quite dark, needs multiple bounces to reach the only (directional) light.
- TeapotMetal (64spp) Simple scene with specular materials
- RingsPhotonMap (128spp) Several caustics, uses Photon Mapping for rendering (all other scenes are using standard path tracing).
- SponzaKilleroo (512spp) Same as SponzaDir, but with an additional golden object that cause caustics and many spikes.
- SibenikCar (256spp) Highly specular materials (car and ground).
- SibenikCrash (256spp) Same as SibenikCar, but with motion blur on the car.

All images have resolution 1024×1024

SanMiguelDof — Simple PBRT rendering (256 spp)

Time: 2732 s

ssim: 0.5913 — relMSE: 2065×10^{-4}

SanMiguelDof — ARNLMF

Time (adaptive sampling and reconstruction): 3313 s

ssim: 0.9318 — relMSE: 90×10^{-4}

SanMiguelDof — RDFCI

Time (rendering + filtering): 2803 s

ssim: 0.9639 — relMSE: 70×10^{-4}

SanMiguelDof — WLR

Time (adaptive sampling and reconstruction): 3506 s

ssim: 0.9585 — relMSE: 68×10^{-4}

SanMiguelDof — LBF

Time (rendering + filtering): 3397 s

ssim: 0.9429 — relMSE: 95×10^{-4}

SanMiguelDof — RHF

Time (rendering + filtering): $2732 + 232 = 2964$ s

ssim: 0.9324 — relMSE: 122×10^{-4}

SanMiguelDof — ours

Time (rendering + filtering): $2732 + 37 = 2769$ s

ssim: 0.9350 — relMSE: 91×10^{-4}

SanMiguelDof — Groundtruth

65536 spp

ssim: 1 — relMSE: 0

SponzaDir — Simple PBRT rendering (256 spp)

Time: 966 s

ssim: 0.8191 — relMSE: 1631×10^{-4}

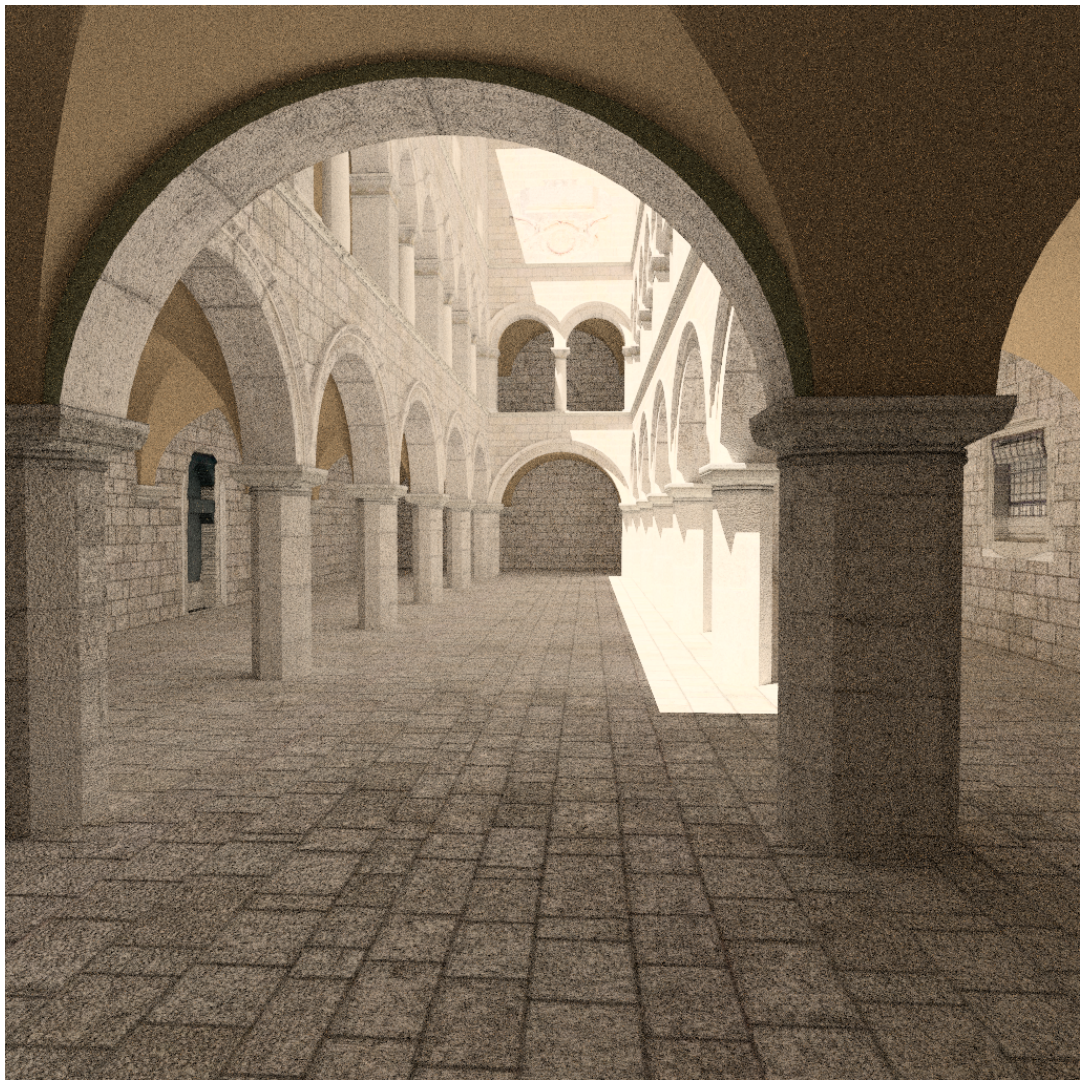

SponzaDir — ARNLMF

Time (adaptive sampling and reconstruction): 1114 s

ssim: 0.9386 — relMSE: 398×10^{-4}

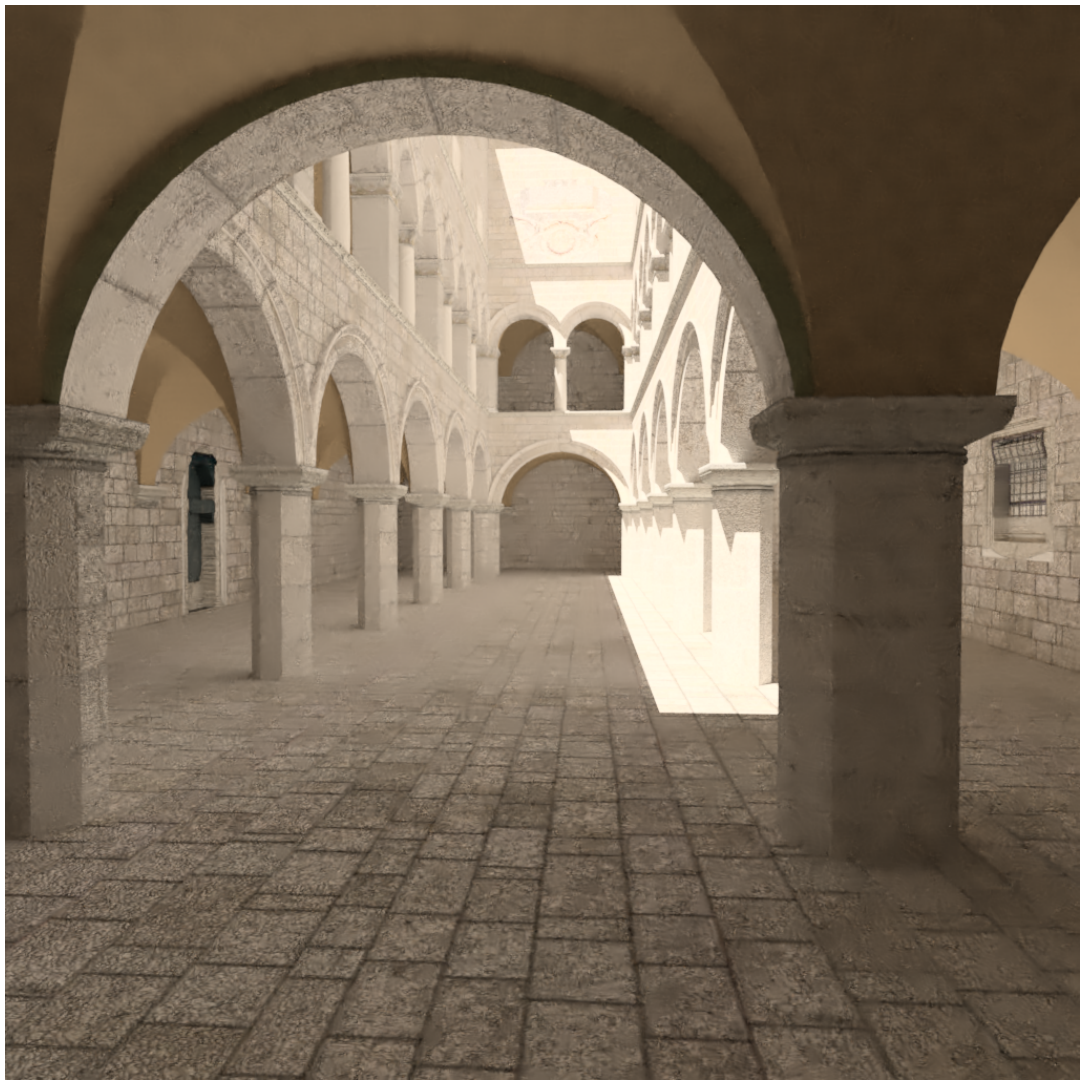

SponzaDir — RDFCI

Time (rendering + filtering): 964 s

ssim: 0.9728 — relMSE: 157×10^{-4}

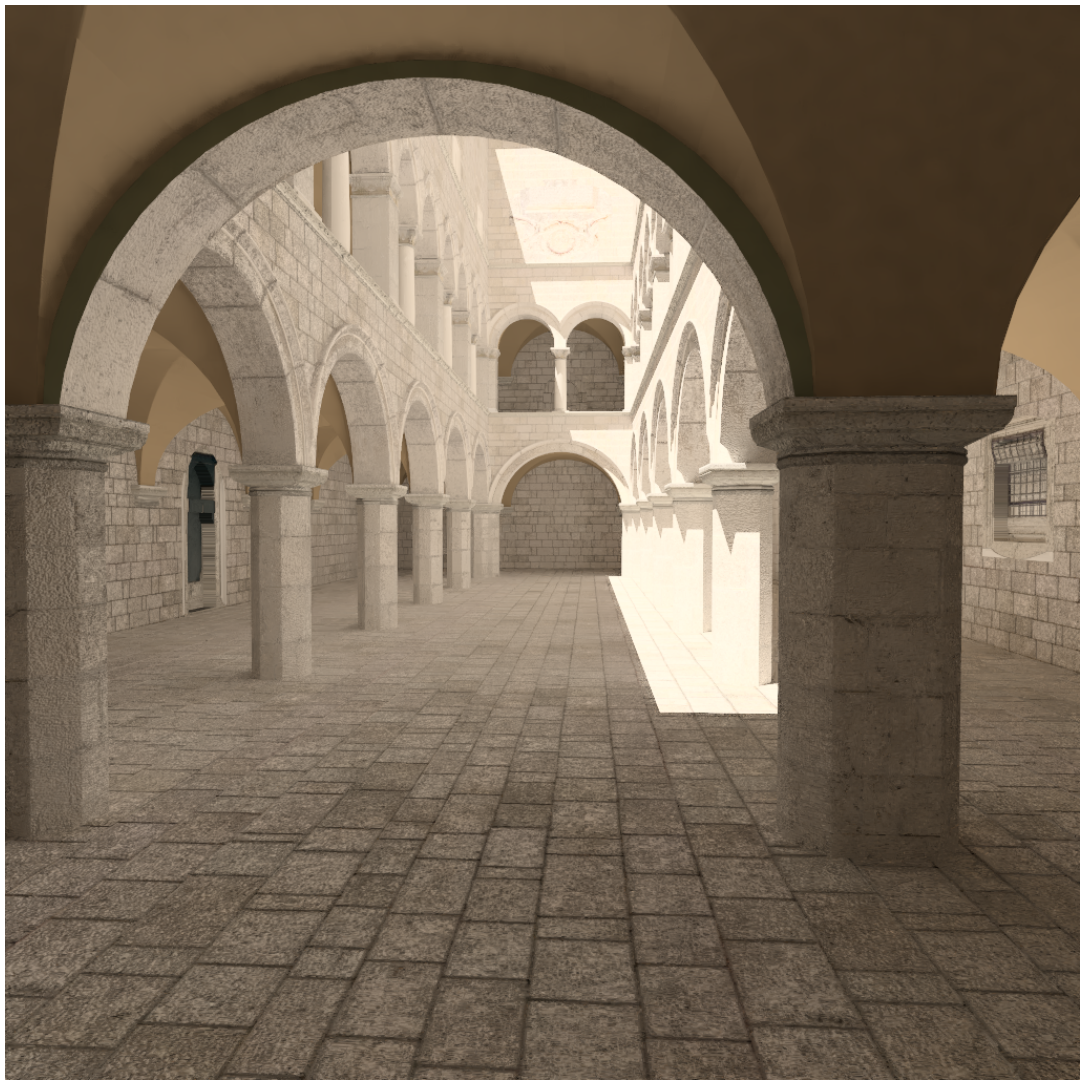

SponzaDir — WLR

Time (adaptive sampling and reconstruction): 1175 s

ssim: 0.9707 — relMSE: 175×10^{-4}

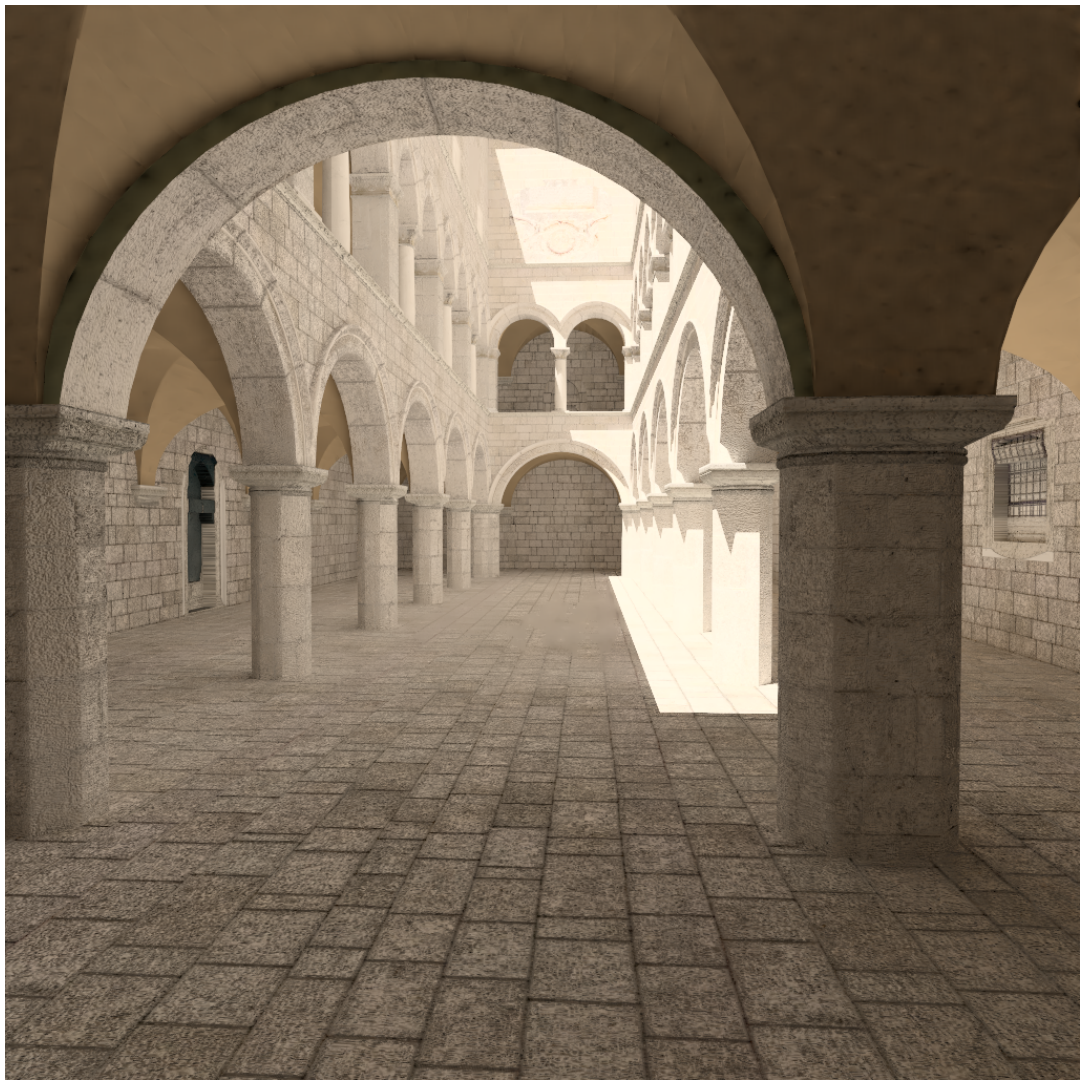

SponzaDir — LBF

Time (rendering + filtering): 1286 s

ssim: 0.9537 — relMSE: 269×10^{-4}

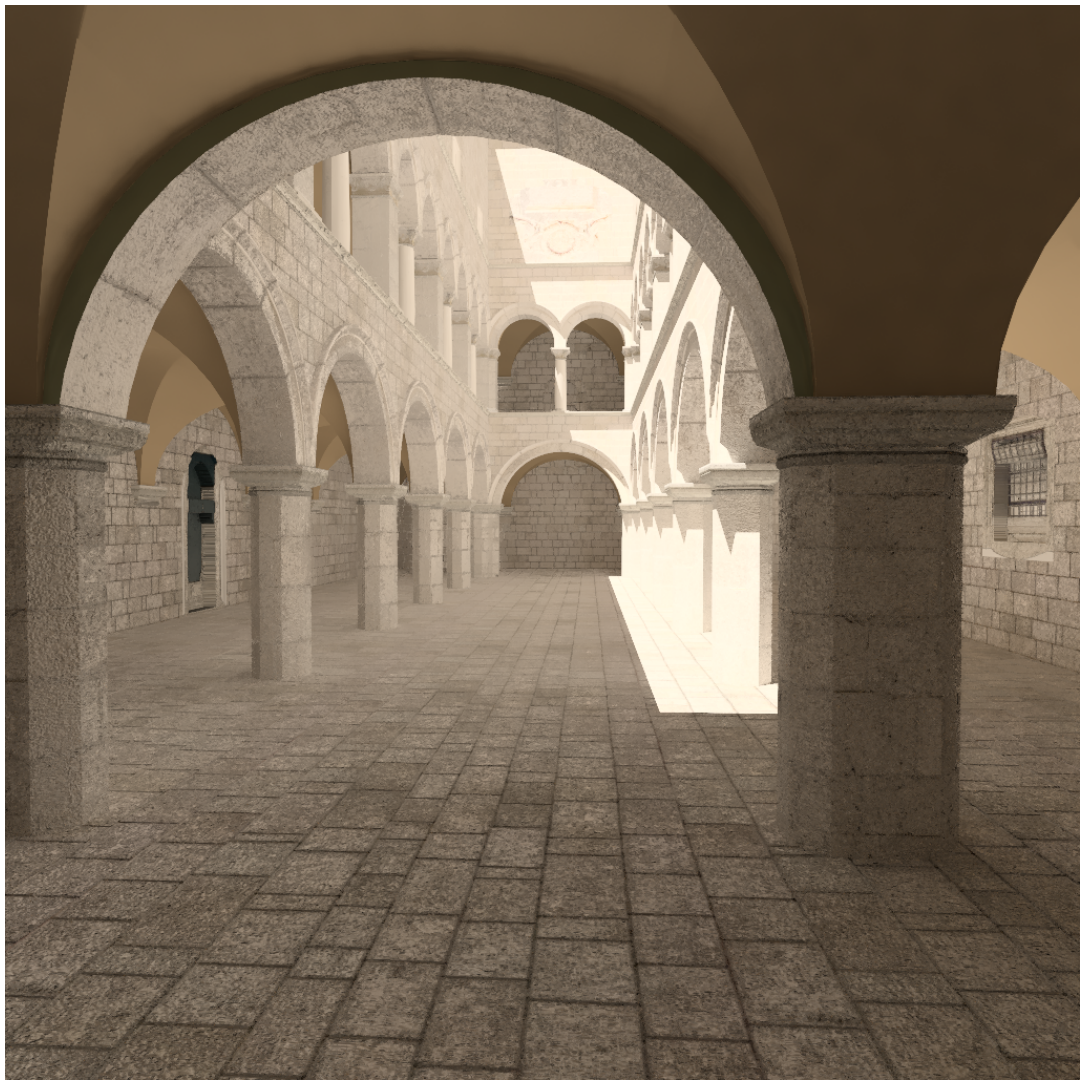

SponzaDir — RHF

Time (rendering + filtering): $966 + 270 = 1236$ s

ssim: 0.9000 — relMSE: 587×10^{-4}

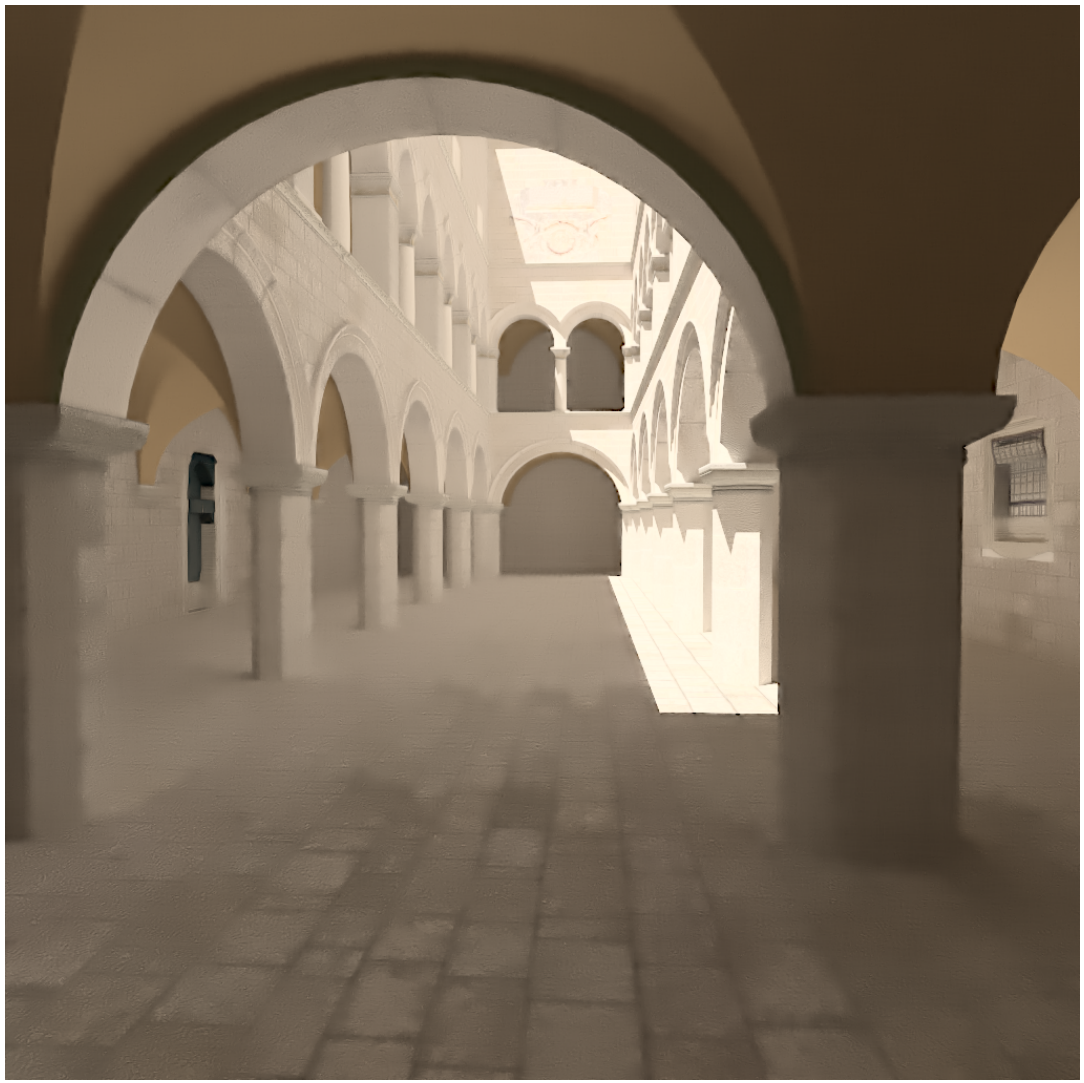

SponzaDir — ours

Time (rendering + filtering): $966 + 22 = 988$ s

ssim: 0.9426 — relMSE: 276×10^{-4}

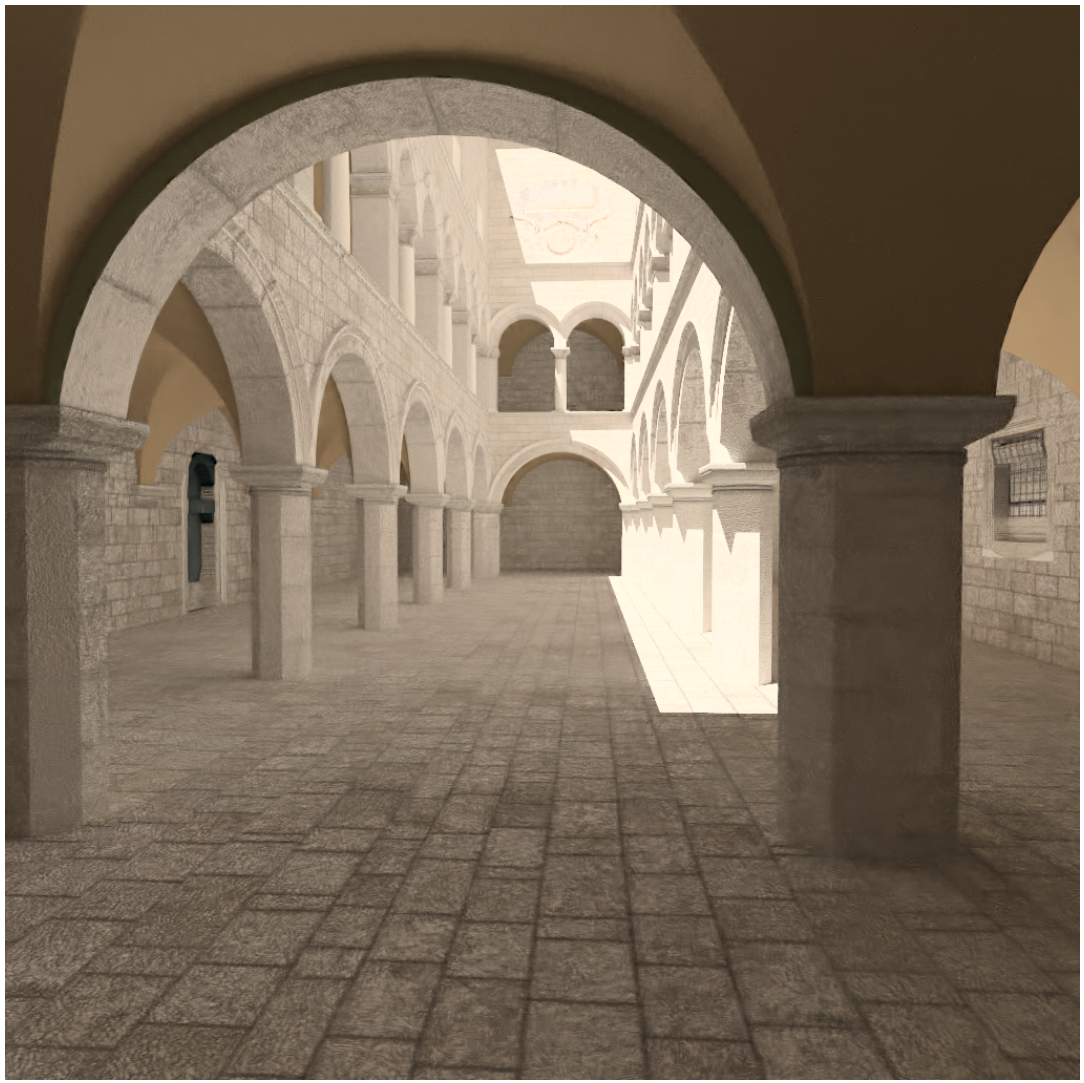

SponzaDir — Groundtruth

65536 spp

ssim: 1 — relMSE: 0

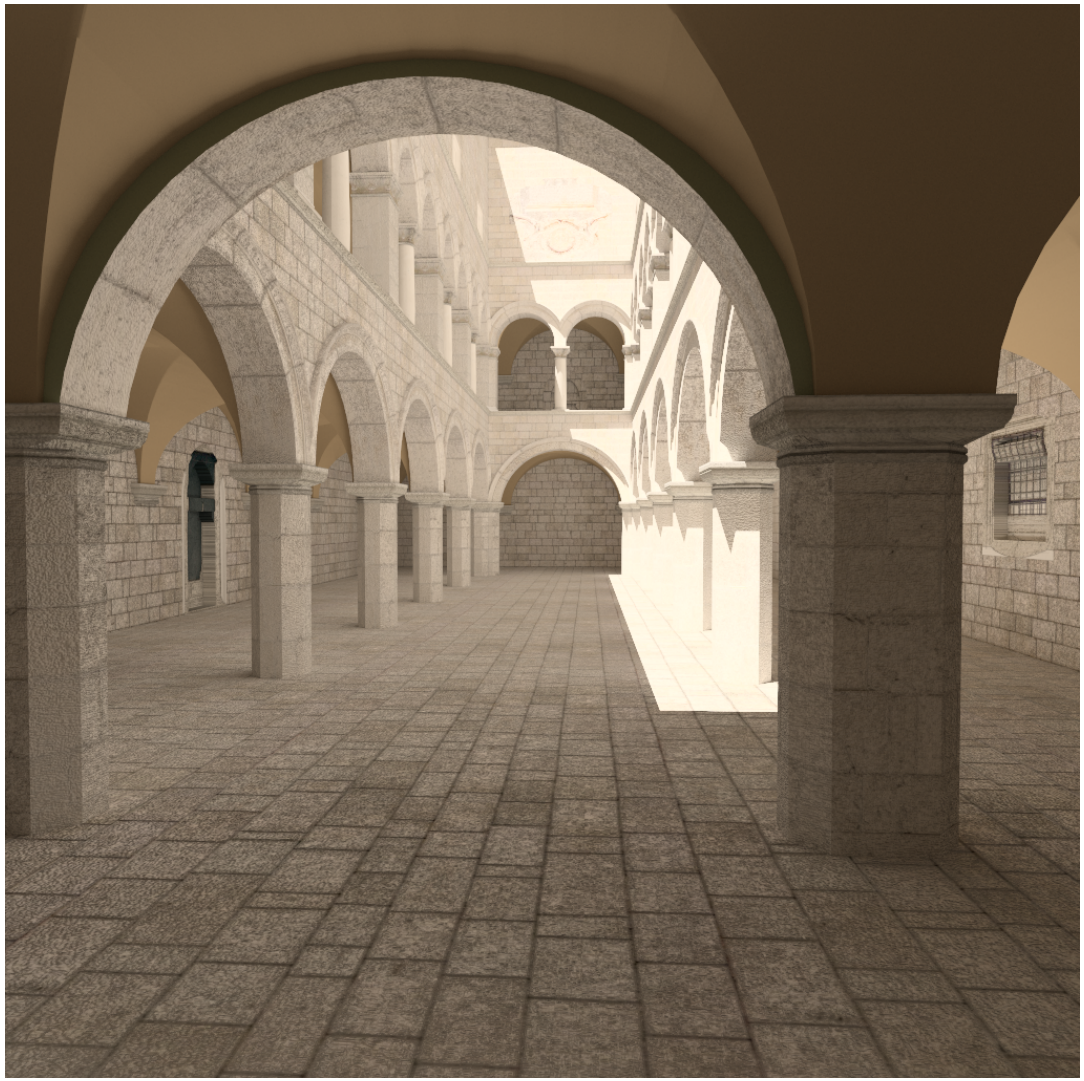

TeapotMetal — Simple PBRT rendering (64 spp)

Time: 186 s

ssim: 0.6422 — relMSE: 1802×10^{-4}

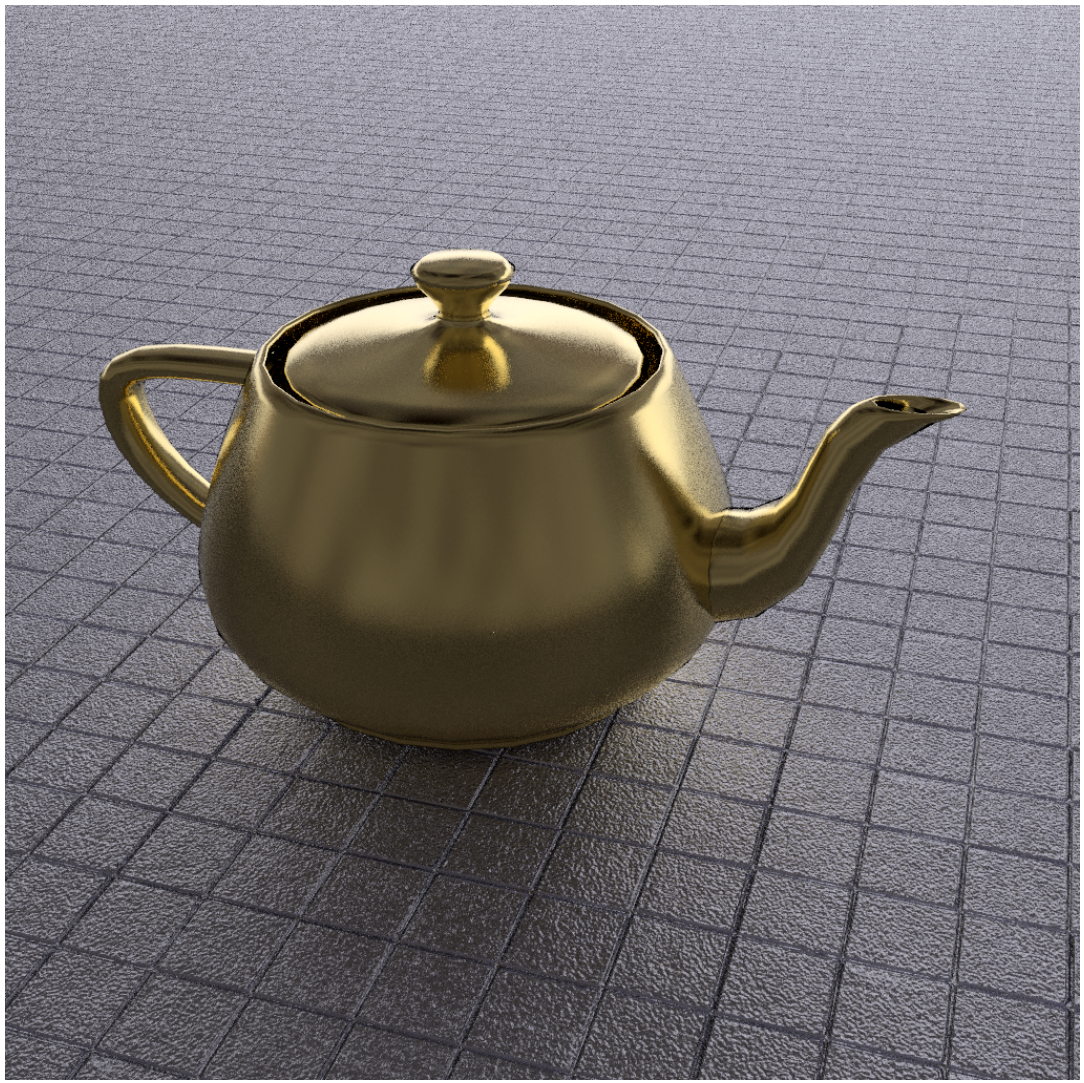

TeapotMetal — ARNLMF

Time (adaptive sampling and reconstruction): 222 s

ssim: 0.7579 — relMSE: 1488×10^{-4}

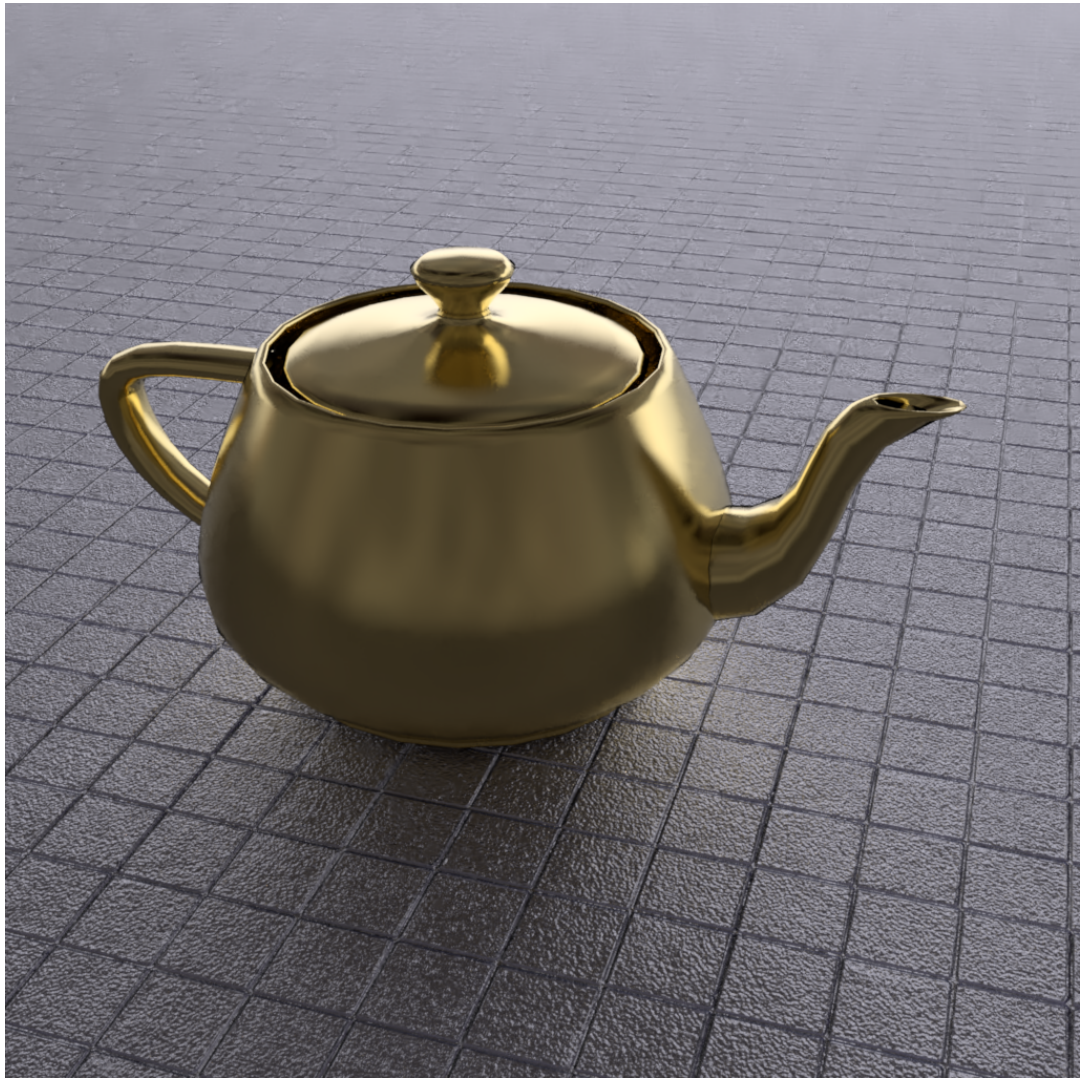

TeapotMetal — RDFCI

Time (rendering + filtering): 165 s

ssim: 0.8671 — relMSE: 641×10^{-4}

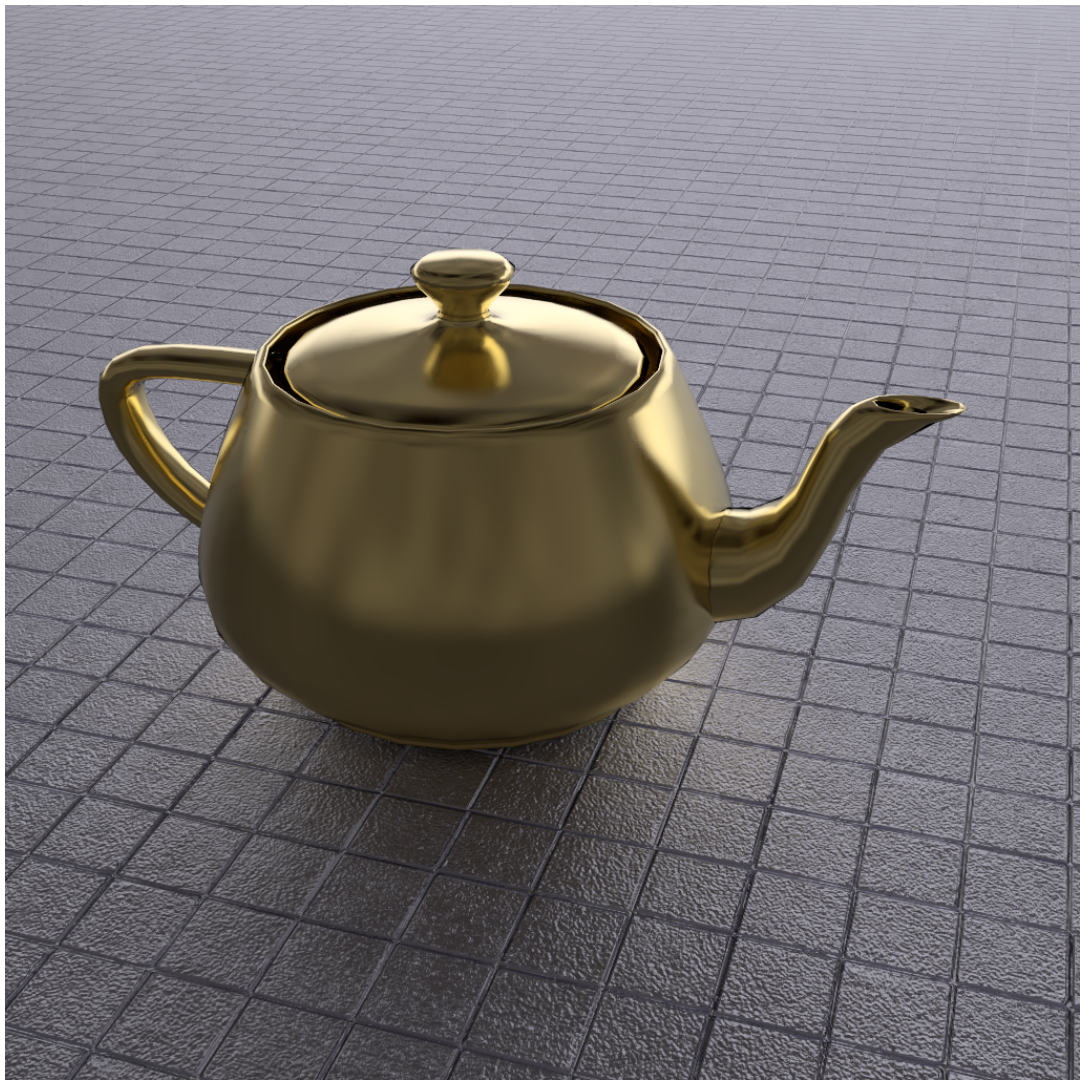

TeapotMetal — WLR

Time (adaptive sampling and reconstruction): 255 s

ssim: 0.8431 — relMSE: 795×10^{-4}

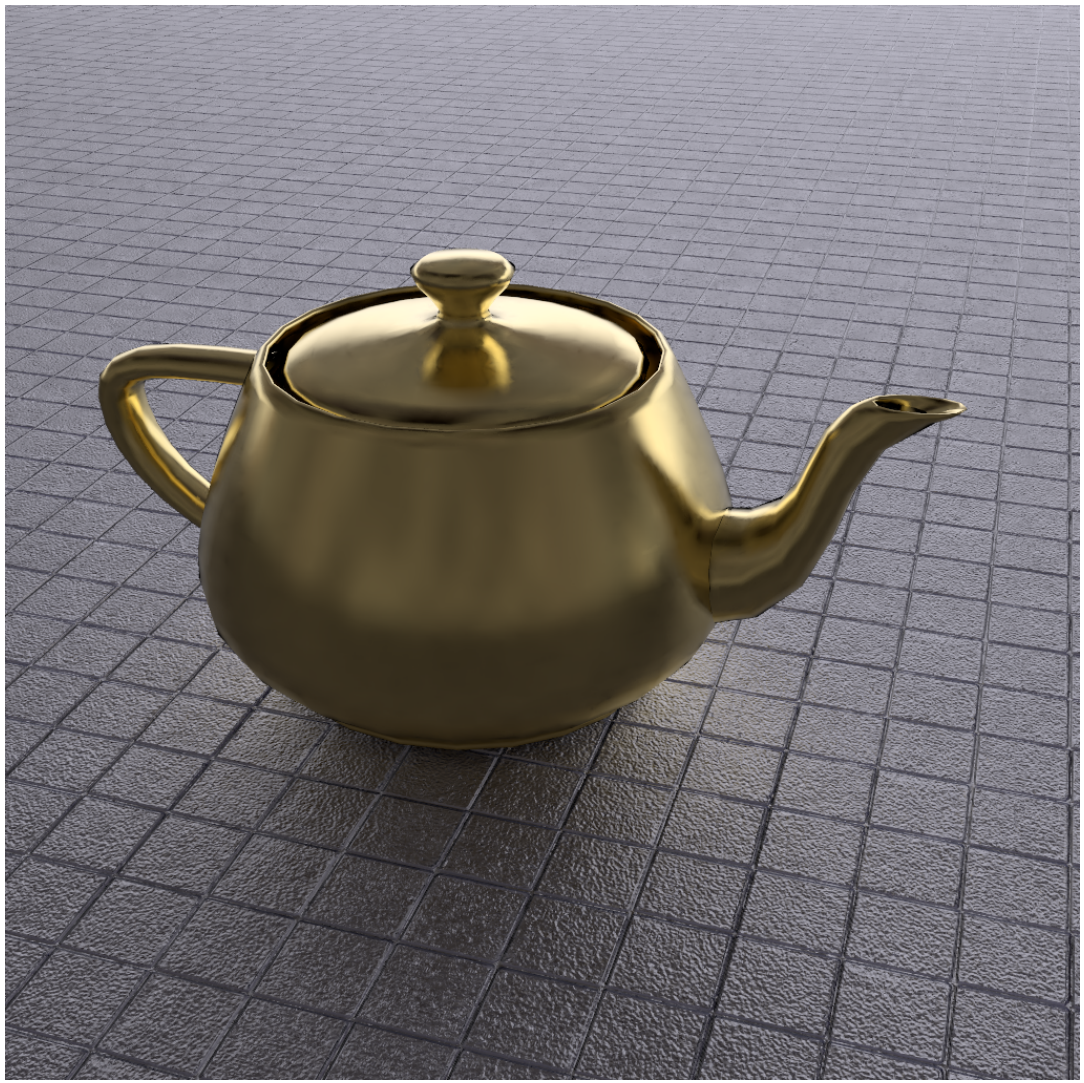

TeapotMetal — LBF

Time (rendering + filtering): 247 s

ssim: 0.8232 — relMSE: 764×10^{-4}

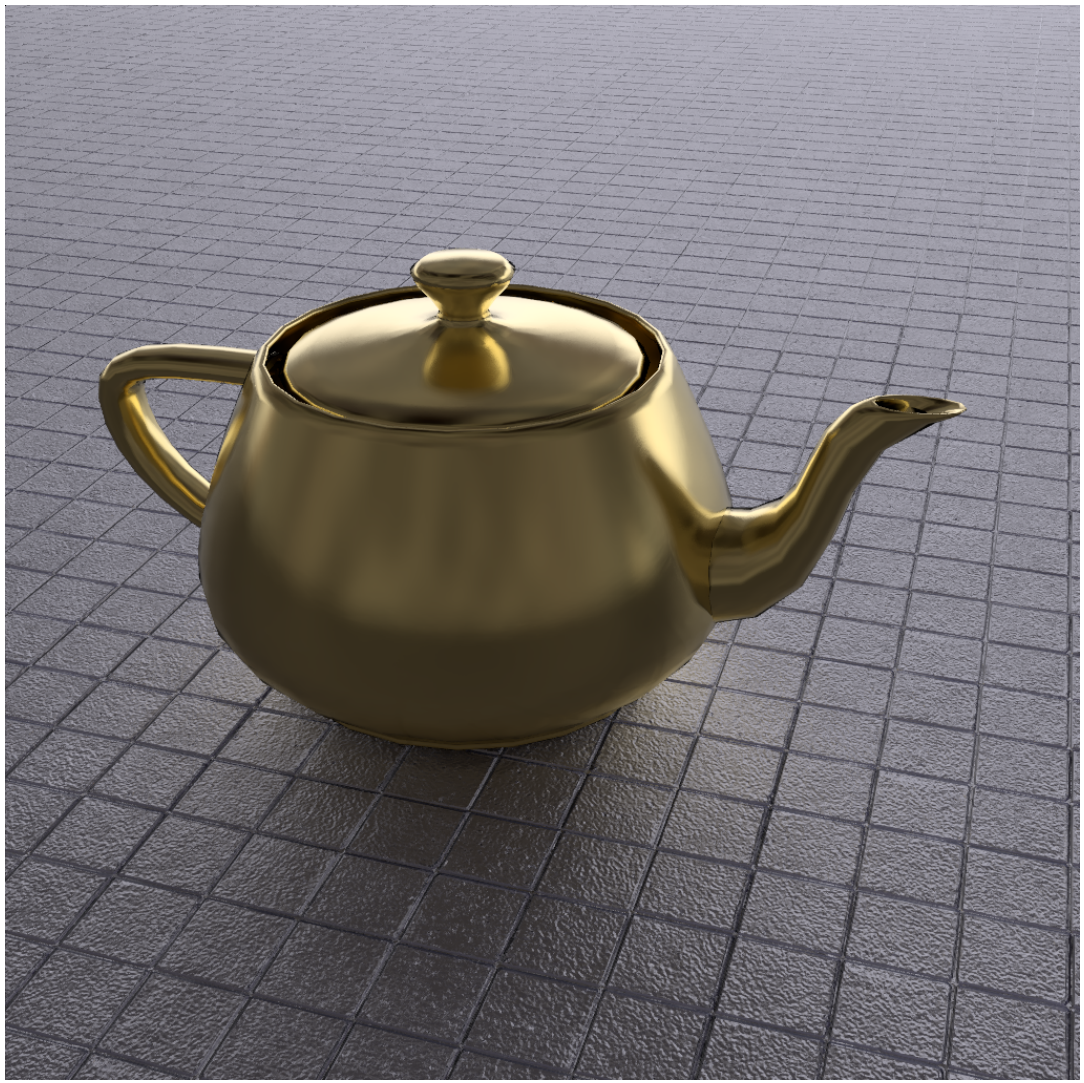

TeapotMetal — RHF

Time (rendering + filtering): $186 + 259 = 445$ s

ssim: 0.7763 — relMSE: 1164×10^{-4}

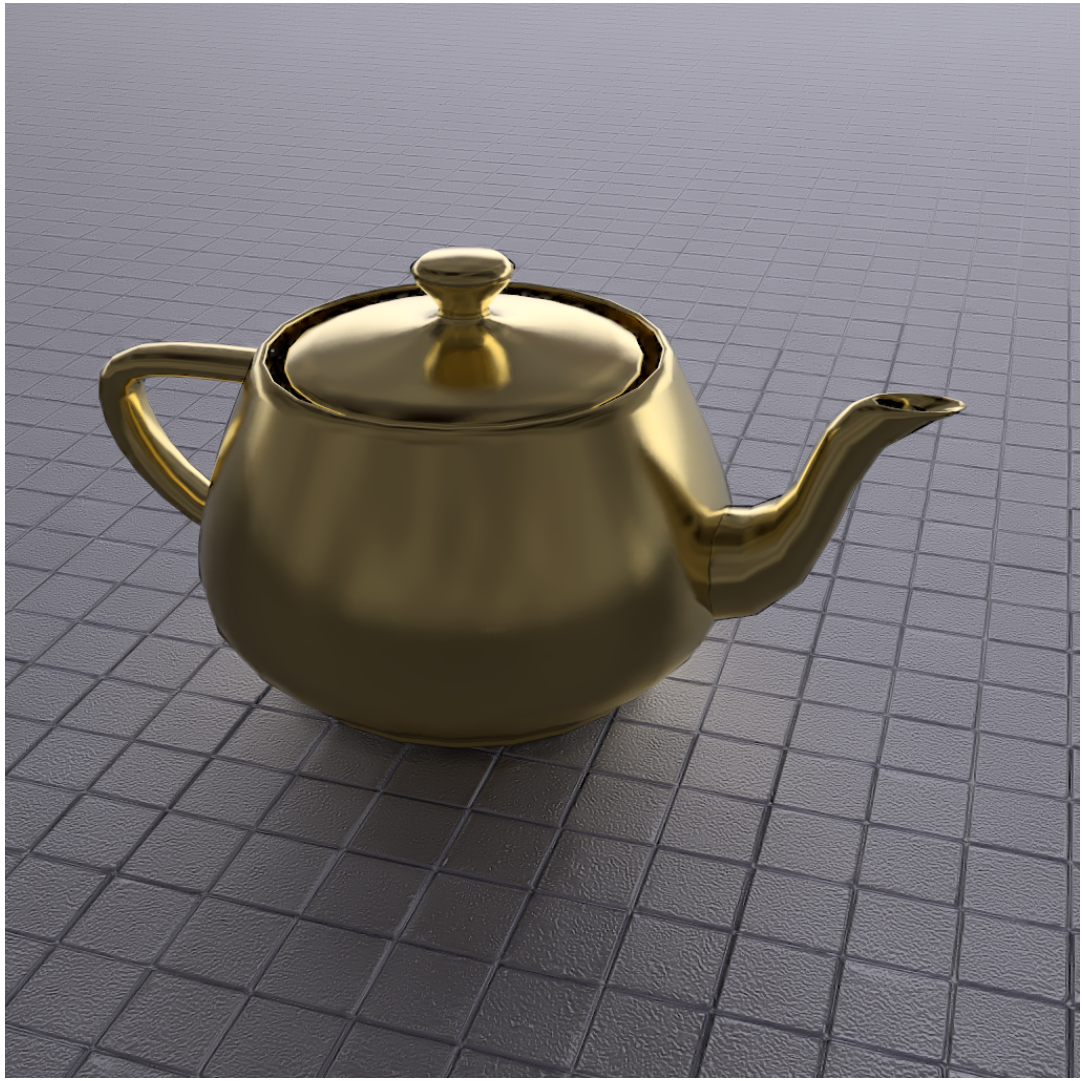

TeapotMetal — ours

Time (rendering + filtering): $186 + 41 = 227$ s

ssim: 0.8160 — relMSE: 841×10^{-4}

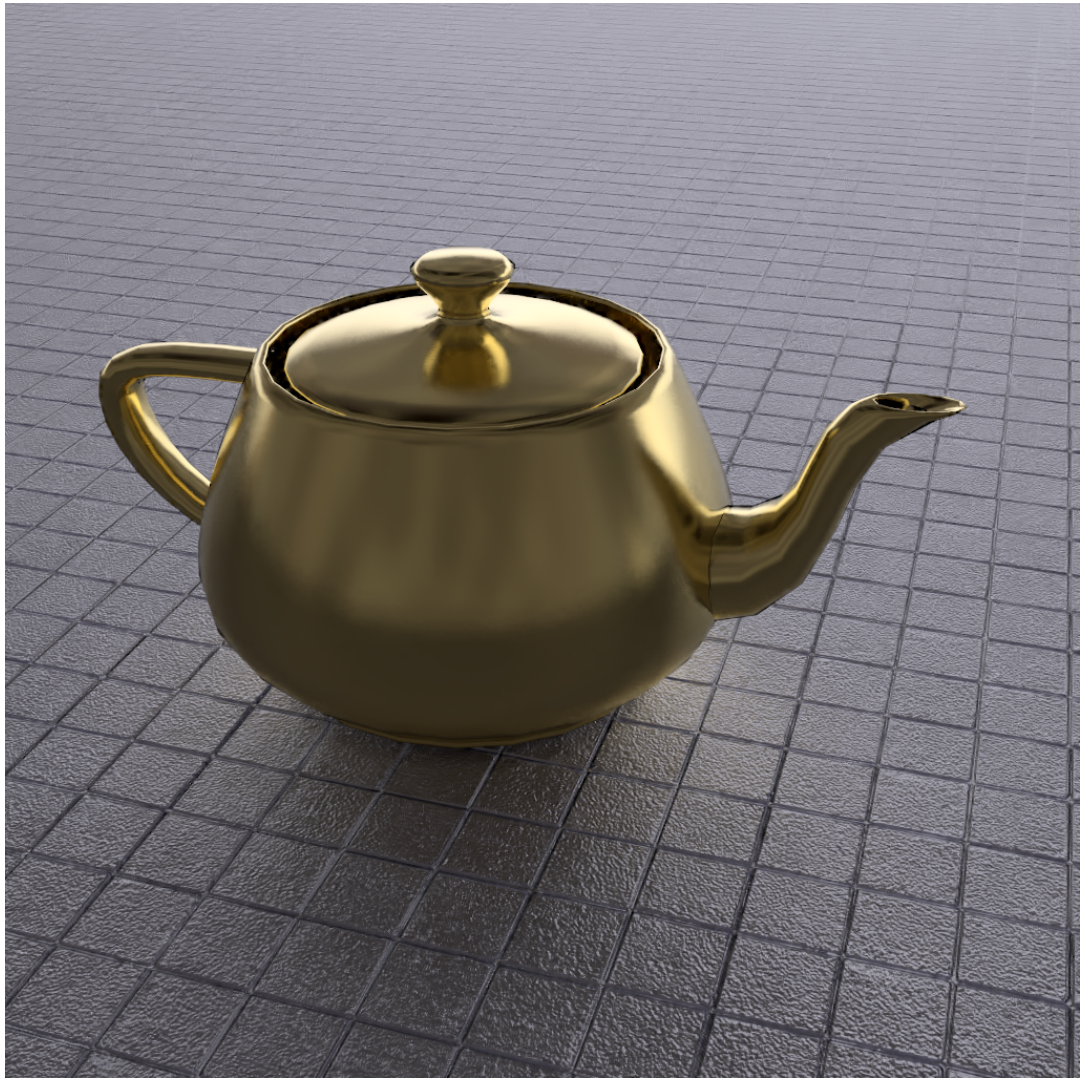

TeapotMetal — Groundtruth

65536 spp

ssim: 1 — relMSE: 0

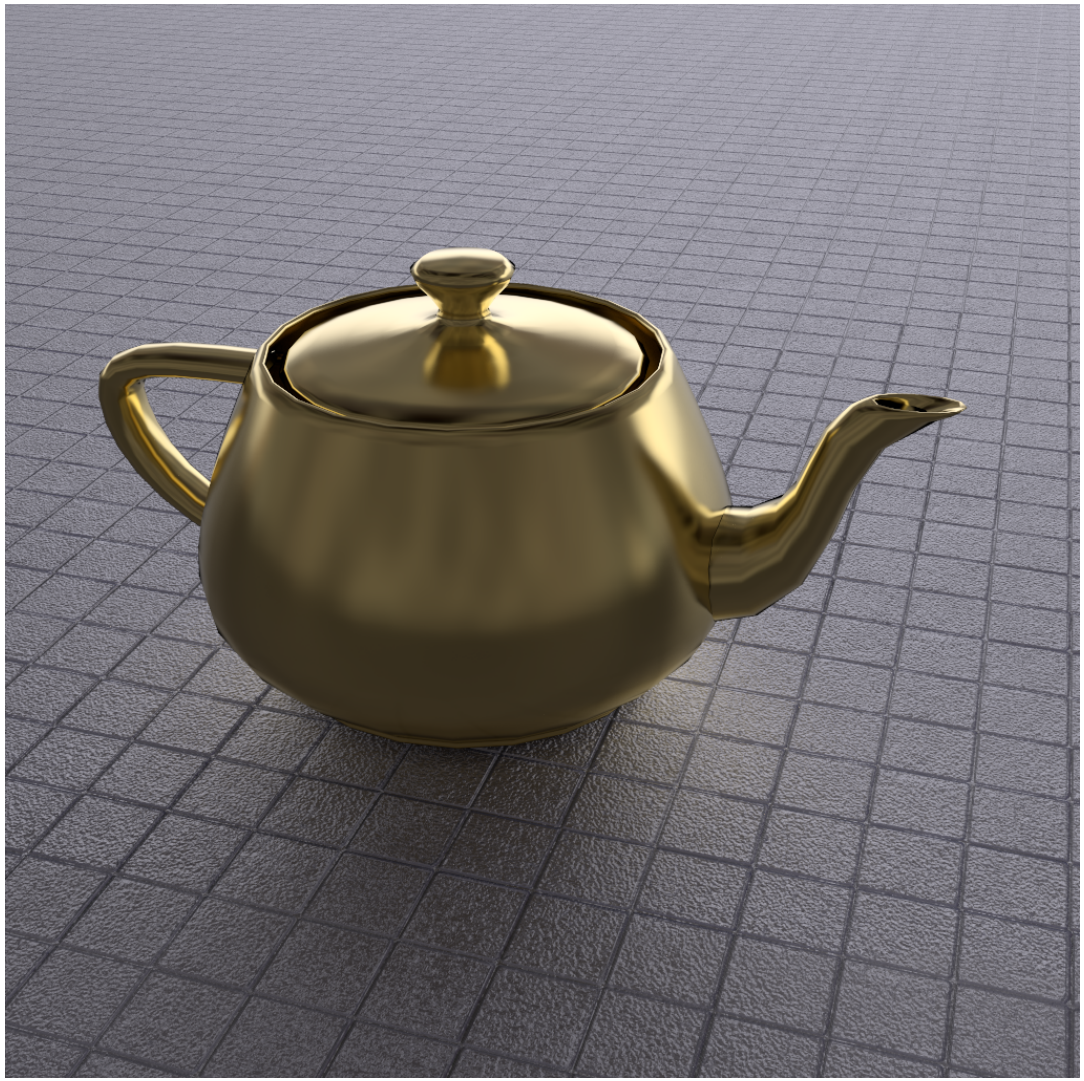

RingsPhotonMap — Simple PBRT rendering (128 spp)

Time: 16904 s

ssim: 0.5358 — relMSE: 1413×10^{-4}

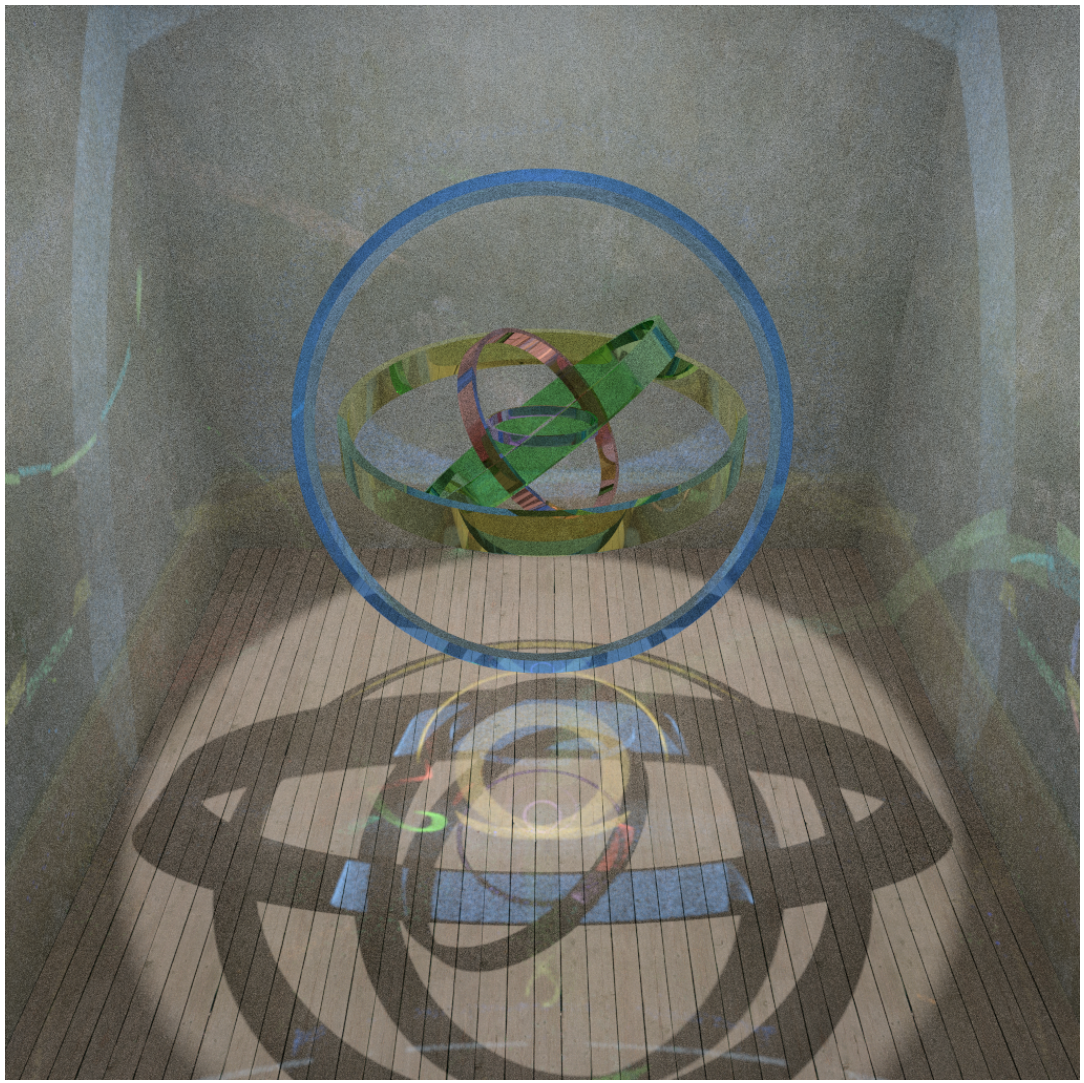

RingsPhotonMap — ARNLMF

Time (adaptive sampling and reconstruction): 16918 s

ssim: 0.8271 — relMSE: 1148×10^{-4}

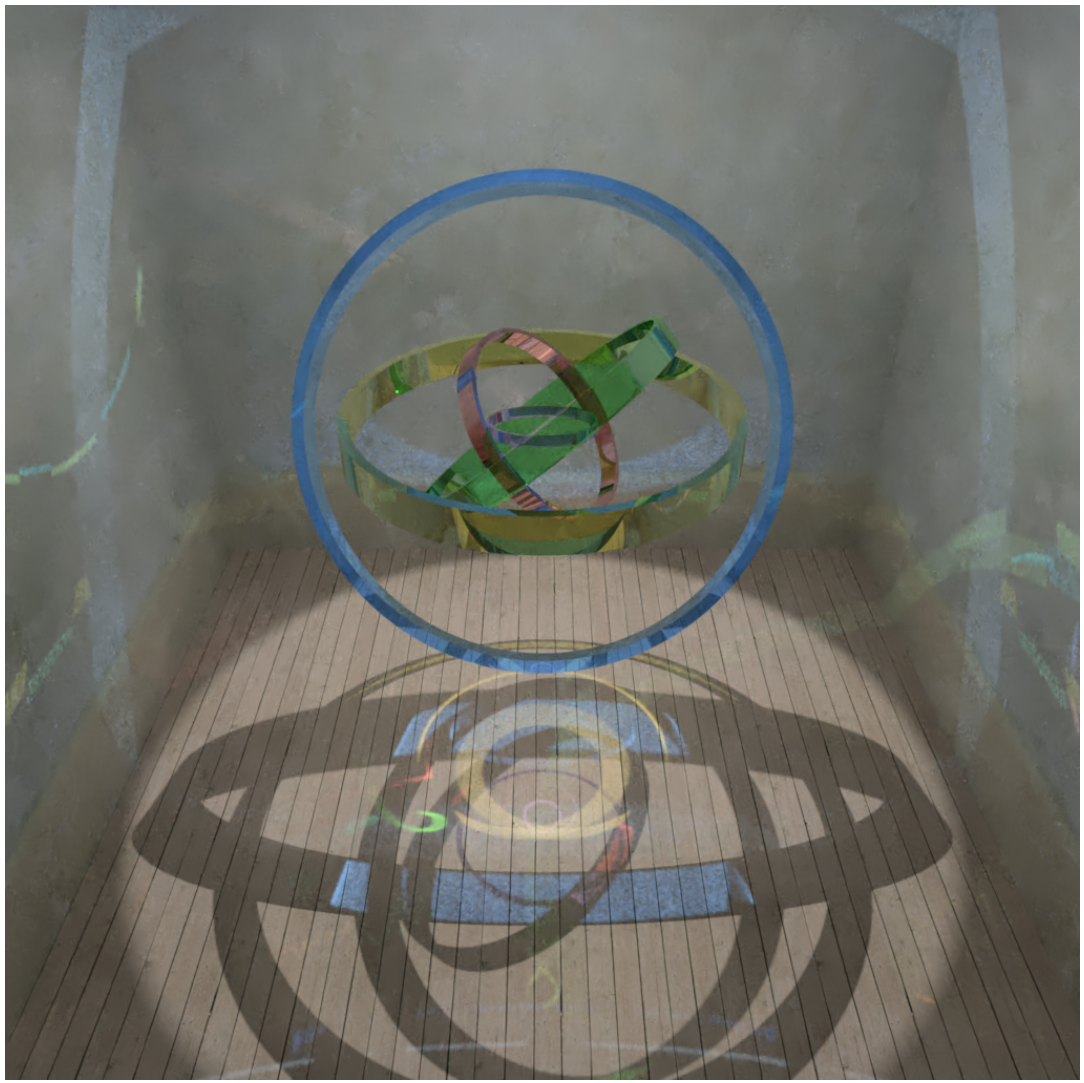

RingsPhotonMap — RDFCI

Time (rendering + filtering): 11084 s

ssim: 0.8106 — relMSE: 1214×10^{-4}

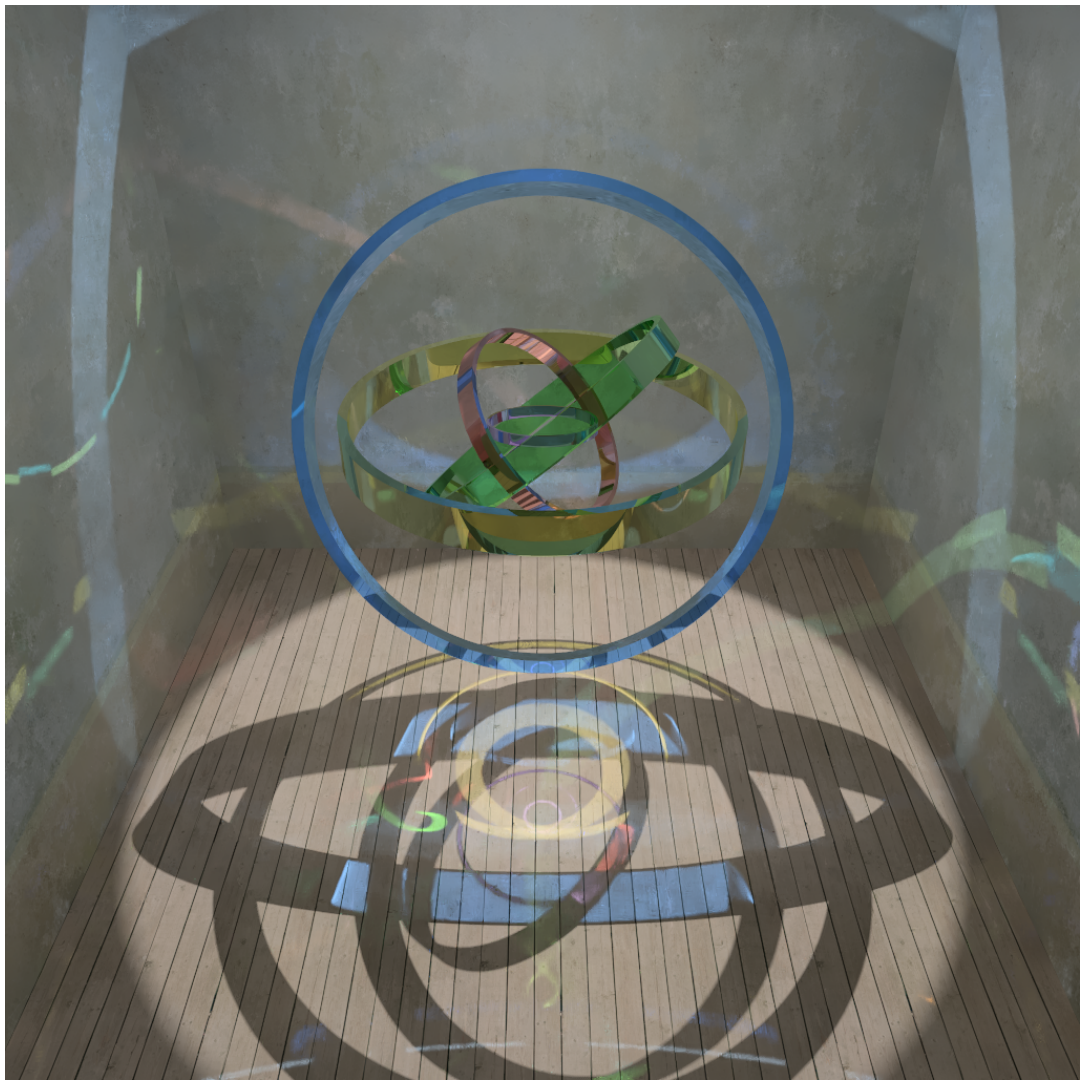

RingsPhotonMap — WLR

Time (adaptive sampling and reconstruction): 14433 s

ssim: 0.8938 — relMSE: 925×10^{-4}

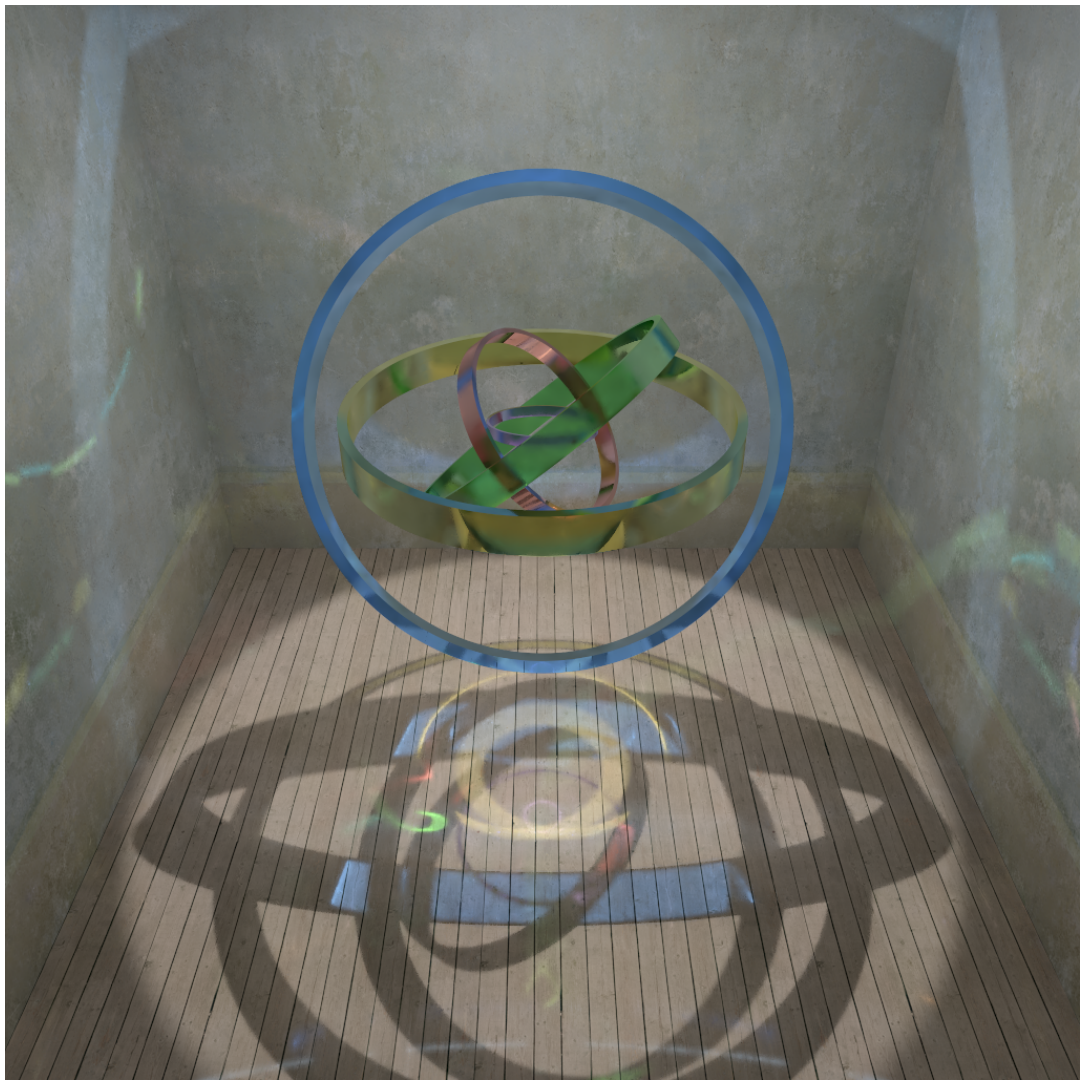

RingsPhotonMap — LBF

Time (rendering + filtering): 17758 s

ssim: 0.7718 — relMSE: 2107×10^{-4}

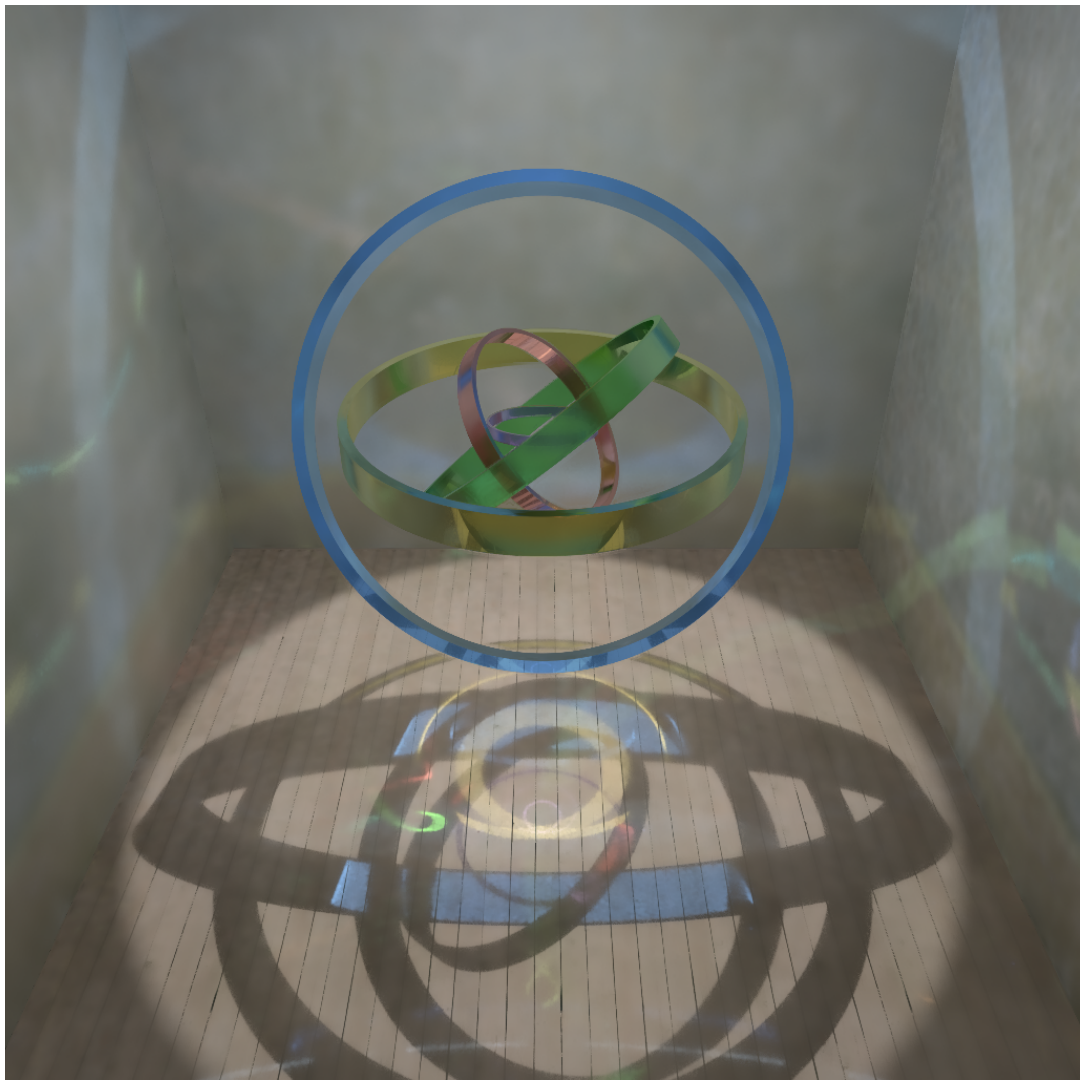

RingsPhotonMap — RHF

Time (rendering + filtering): $16904 + 286 = 17190$ s

ssim: 0.8460 — relMSE: 886×10^{-4}

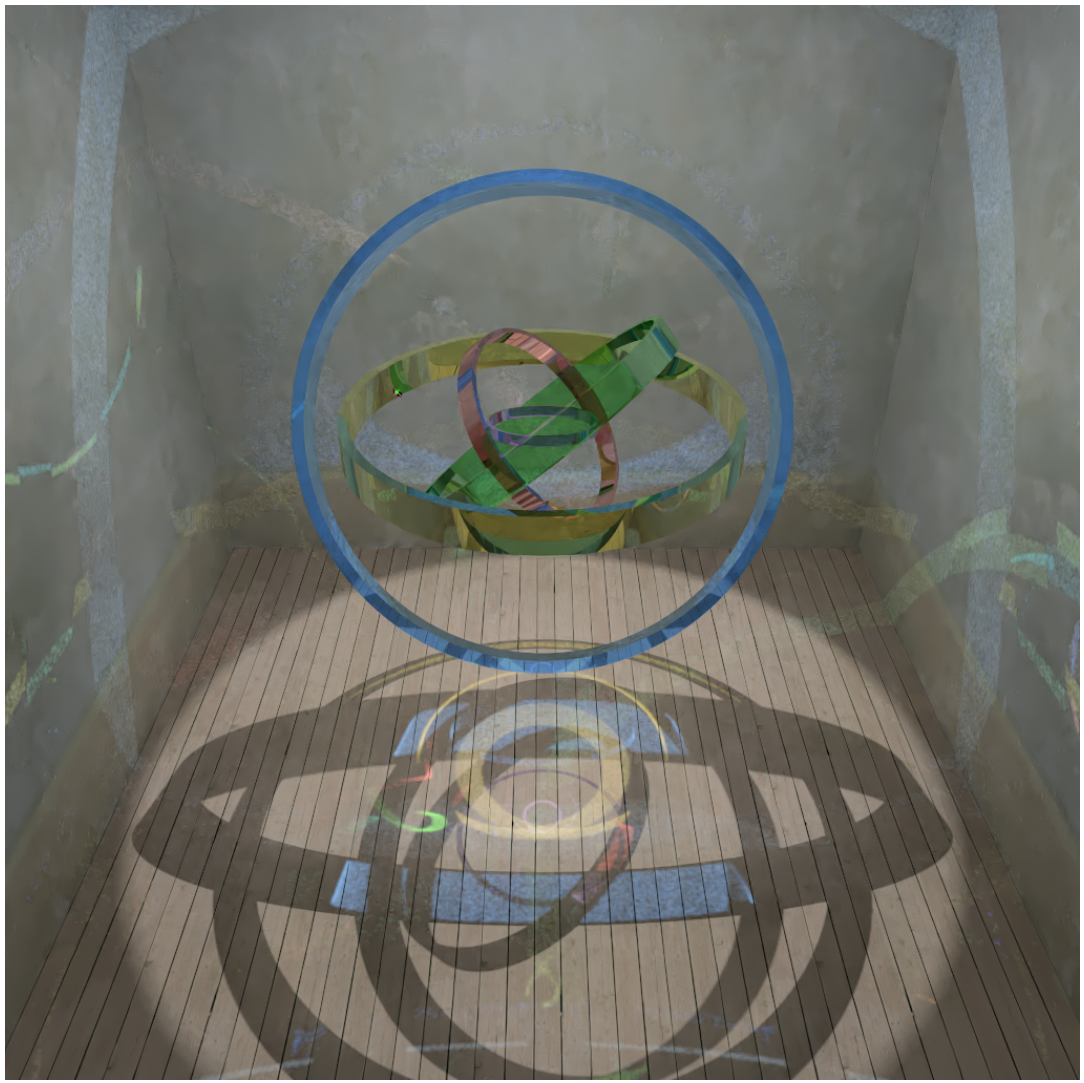

RingsPhotonMap — ours

Time (rendering + filtering): $16904 + 56 = 16960$ s

ssim: 0.8433 — relMSE: 498×10^{-4}

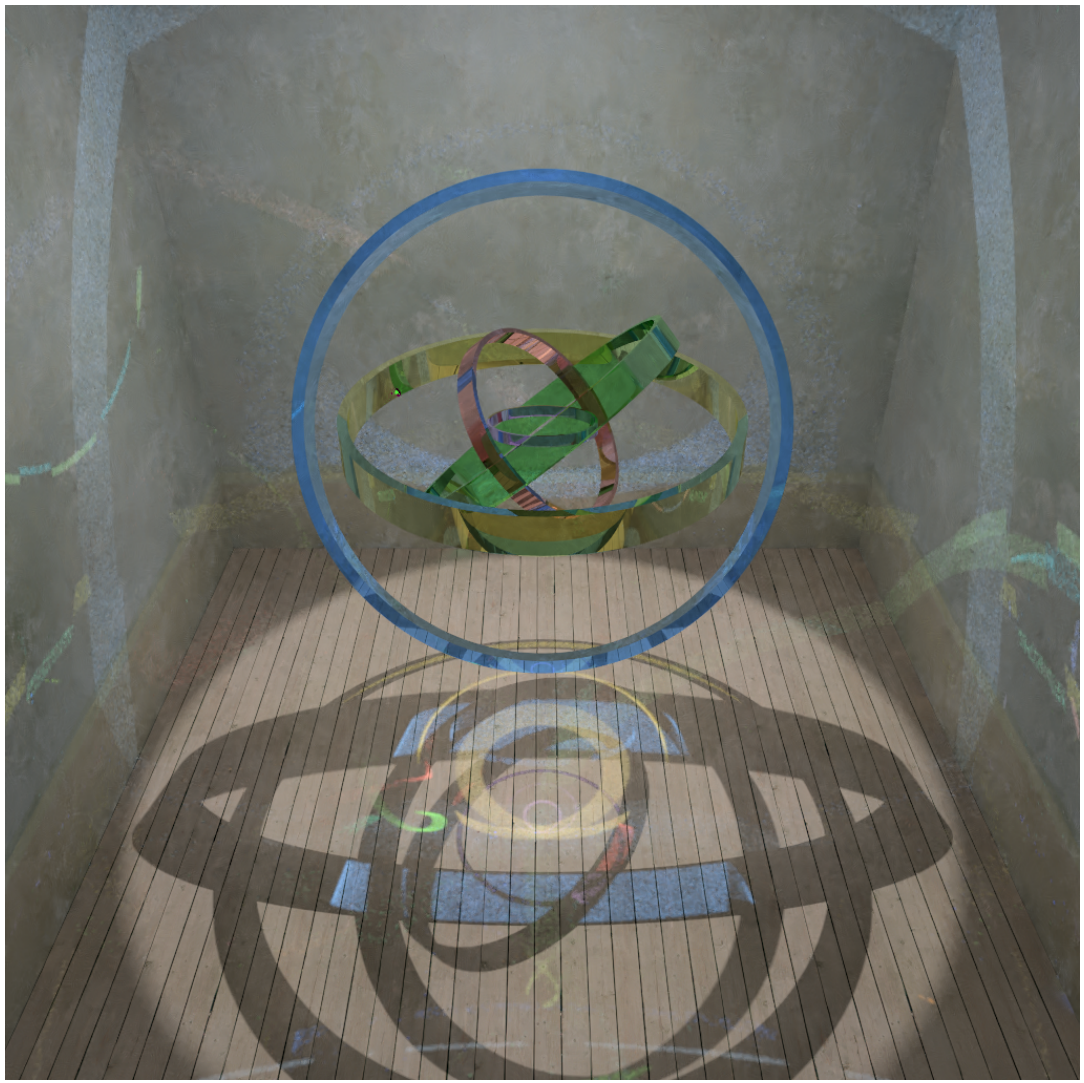

RingsPhotonMap — Groundtruth

4096 spp

ssim: 1 — relMSE: 0

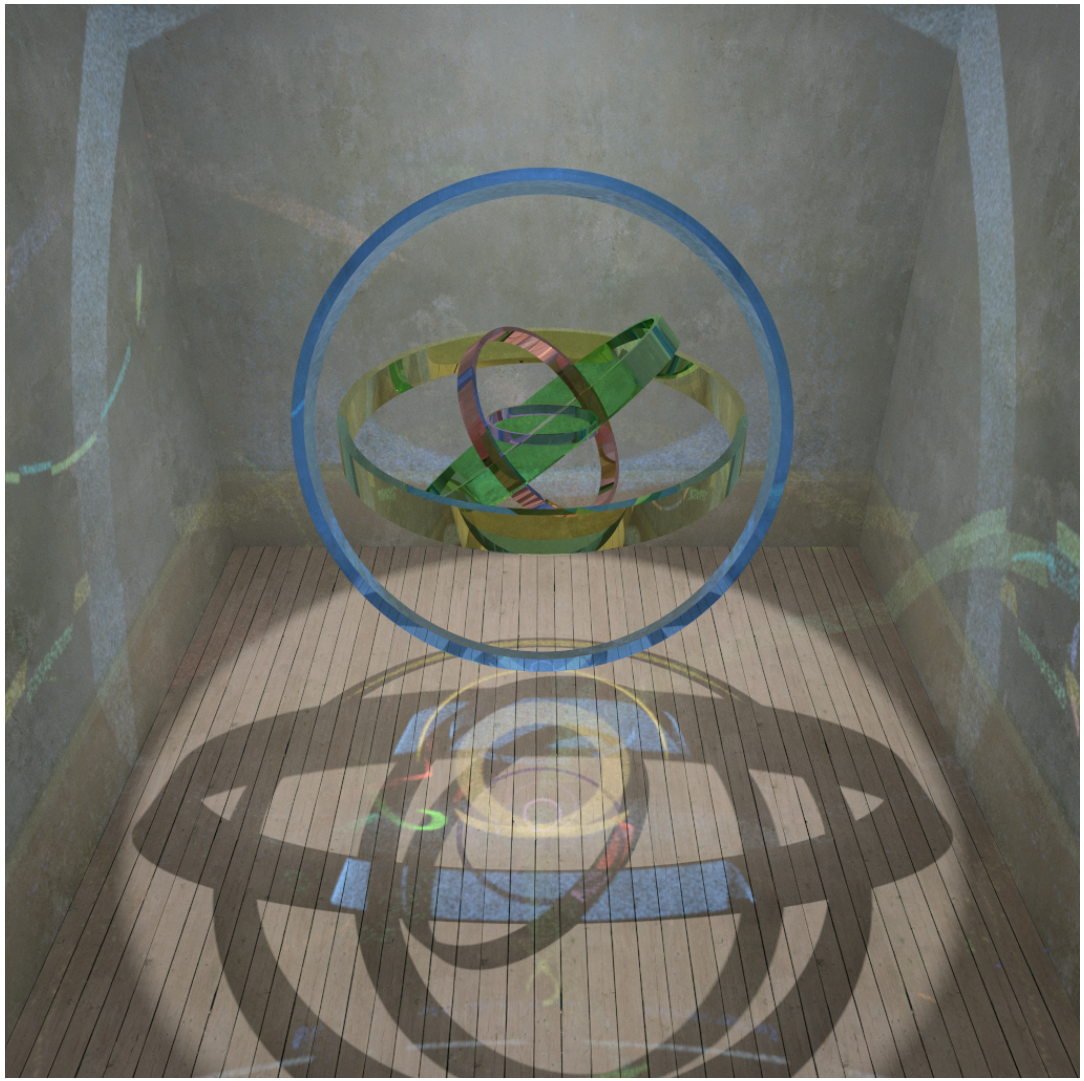

SponzaKilleroo — Simple PBRT rendering (512 spp)

Time: 2109 s

ssim: 0.7830 — relMSE: 7125×10^{-4}

SponzaKilleroo — ARNLMF

Time (adaptive sampling and reconstruction): 2373 s

ssim: 0.9459 — relMSE: 430×10^{-4}

SponzaKilleroo — RDFCI

Time (rendering + filtering): 2150 s

ssim: 0.9652 — relMSE: 245×10^{-4}

SponzaKilleroo — WLR

Time (adaptive sampling and reconstruction): 2374 s

ssim: 0.9671 — relMSE: 260×10^{-4}

SponzaKilleroo — LBF

Time (rendering + filtering): 2735 s

ssim: 0.9572 — relMSE: 257×10^{-4}

SponzaKilleroo — RHF

Time (rendering + filtering): $2109 + 276 = 2385$ s

ssim: 0.9300 — relMSE: 489×10^{-4}

SponzaKilleroo — ours

Time (rendering + filtering): $2109 + 28 = 2137$ s

ssim: 0.9520 — relMSE: 287×10^{-4}

SponzaKilleroo — ours with spike-removal prefiltering

Time (rendering + filtering): $2109 + 28 = 2137$ s

ssim: 0.9523 — relMSE: 267×10^{-4}

SponzaKilleroo — Groundtruth

65536 spp

ssim: 1 — relMSE: 0

SibenikCar — Simple PBRT rendering (256 spp)

Time: 2495 s

ssim: 0.6338 — relMSE: 995×10^{-4}

SibenikCar — ARNLMF

Time (adaptive sampling and reconstruction): 2703 s

ssim: 0.9541 — relMSE: 529×10^{-4}

SibenikCar — RDFCI

Time (rendering + filtering): 2419 s

ssim: 0.9694 — relMSE: 53×10^{-4}

SibenikCar — WLR

Time (adaptive sampling and reconstruction): 2701 s

ssim: 0.9599 — relMSE: 102×10^{-4}

SibenikCar — LBF

Time (rendering + filtering): 2911 s

ssim: 0.9495 — relMSE: 179×10^{-4}

SibenikCar — RHF

Time (rendering + filtering): $2495 + 258 = 2753$ s

ssim: 0.9637 — relMSE: 363×10^{-4}

SibenikCar — ours

Time (rendering + filtering): $2495 + 33 = 2528$ s

ssim: 0.9602 — relMSE: 95×10^{-4}

SibenikCar — Groundtruth

65536 spp

ssim: 1 — relMSE: 0

SibenikCrash — Simple PBRT rendering (256 spp)

Time: 3925 s

ssim: 0.6093 — relMSE: 1038×10^{-4}

SibenikCrash — ARNLMF

Time (adaptive sampling and reconstruction): 4212 s

ssim: 0.9576 — relMSE: 45×10^{-4}

SibenikCrash — RDFCI

Time (rendering + filtering): 3557 s

ssim: 0.9712 — relMSE: 32×10^{-4}

SibenikCrash — WLR

Time (adaptive sampling and reconstruction): 4082 s
ssim: 0.9656 — relMSE: 46×10^{-4}

SibenikCrash — LBF

Time (rendering + filtering): 4329 s

ssim: 0.9585 — relMSE: 64×10^{-4}

SibenikCrash — RHF

Time (rendering + filtering): $3925 + 255 = 4180$ s

ssim: 0.9662 — relMSE: 38×10^{-4}

SibenikCrash — ours

Time (rendering + filtering): $3925 + 29 = 3954$ s

ssim: 0.9625 — relMSE: 42×10^{-4}

SibenikCrash — Groundtruth

65536 spp

ssim: 1 — relMSE: 0

Comparison diagrams