Eurographics Symposium on Rendering - Experimental Ideas & Implementations (2018)
T. Hachisuka and W. Jakob (Editors)

An Improved Multiple Importance Sampling Heuristic for Density Estimates in Light Transport Simulations Supplemental: Visual Comparison

J. Jendersie and T. Grosch

TU Clausthal, Germany

This document contains additional non-cropped full resolution images for visual convergence comparison. Please zoom in to view the full resolution.

Also, it provides further curves using even more error metrics:

$$\text{RMSE} := \sqrt{\frac{1}{N}\sum_i^N (x_i - r_i)^2},$$

$$\text{RMSLE} := \sqrt{\frac{1}{N}\sum_i^N (\log(1+r_i) - \log(1+x_i))^2},$$

$$\text{SRRMSE} := \sqrt{\frac{1}{N}\sum_i^N \left(\frac{x_i - r_i}{(x_i + r_i)/2}\right)^2}$$

$$\text{SSIM} := \frac{(2\mu_{xi}\mu_{ri} + C_1)(2\sigma_{xri} + C_2)}{(\mu_{xi}^2 + \mu_{ri}^2 + C_1)(\sigma_{xi}^2 + \sigma_{ri}^2 + C_2)},$$

$$\text{and SMAPE} := \frac{1}{N}\sum_i^N 2\left|\frac{r_i - x_i}{r_i + x_i}\right|$$

From those RMSE, SRRMSE and SSIM [WBSS04] are contained in the paper too. Again N is the number of pixels (times the color channels), r_i is the reference value, x_i the current image value and $\mu_{xi}, \mu_{ri}, \sigma_{xi}, \sigma_{ri}, \sigma_{xri}$ are local image statistics from a small neighborhood of pixels.

RMSE is the most commonly used criteria. However, it penalizes errors in bright regions stronger. Thus, it happens that a single firefly or aliasing difference in a bright region dominate the entire error.

RMSLE and SRRMSE both fix this issue by applying the logarithm before comparing the image and by using a relative scaling respectively. Both can be justified by perception. First, our perception is more logarithmic than linear. Second, we have a relative threshold for the *just noticeable difference* (Weber-Fechner Law). Therefore, both correlate stronger with the visual comparison.

SSIM is designed to be the perceived error. By using statistics it penalizes kinds of errors differently. A constant over- or underillumination is not as bad as discontinuity artifacts or noise. Compared to RMSLE and SRRMSE it shows similar characteristic in the comparison below. However, crossover points, at which the quality-order of methods change, move towards higher iterations.

Start Radius for Photon Queries

We have chosen the initial photon query radius manually for each scene to provide the best looking results for any of the VCM methods. Smaller radii would increase the noise and hence variance whereas larger radii would blur stronger. In any case, our change of p_{acc} is still meaningful. Invalid decisions, as in the teaser or the double merge scene, will also appear for different radii.

How To Use this Document

This document uses interactive layers to make pixel-by-pixel comparisons of the methods possible. According to the documentation of the LaTeX-OCGX package, this is compatible with *Acrobat Reader*, *Foxit Reader*, *PDF-XChange-Viewer* and *Evince*. The built-in Windows-*Reader*, *SumatraPDF* and the usual Browser-Plugins do not work. For other viewers it might be possible to manually show/hide layers, if the buttons do not work. Also, to switch between layers if zoomed in and if buttons are outside the current view, please use the viewer's layer switching capabilities.

References

[PJH17] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, 3 ed. Morgan Kaufmann, 2017. URL: http://www.pbrt.org/. 4, 5, 7, 10

[Vev18] VEVODA P.: Small UPBP Repository, 2018. URL: https://github.com/PetrVevoda/smallupbp/tree/master/scenes/mirrorballs. 9

[WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image Quality Assessment: from Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing (TIP) 13*, 4 (2004), 600–612. URL: https://doi.org/10.1109/TIP.2003.819861. 1

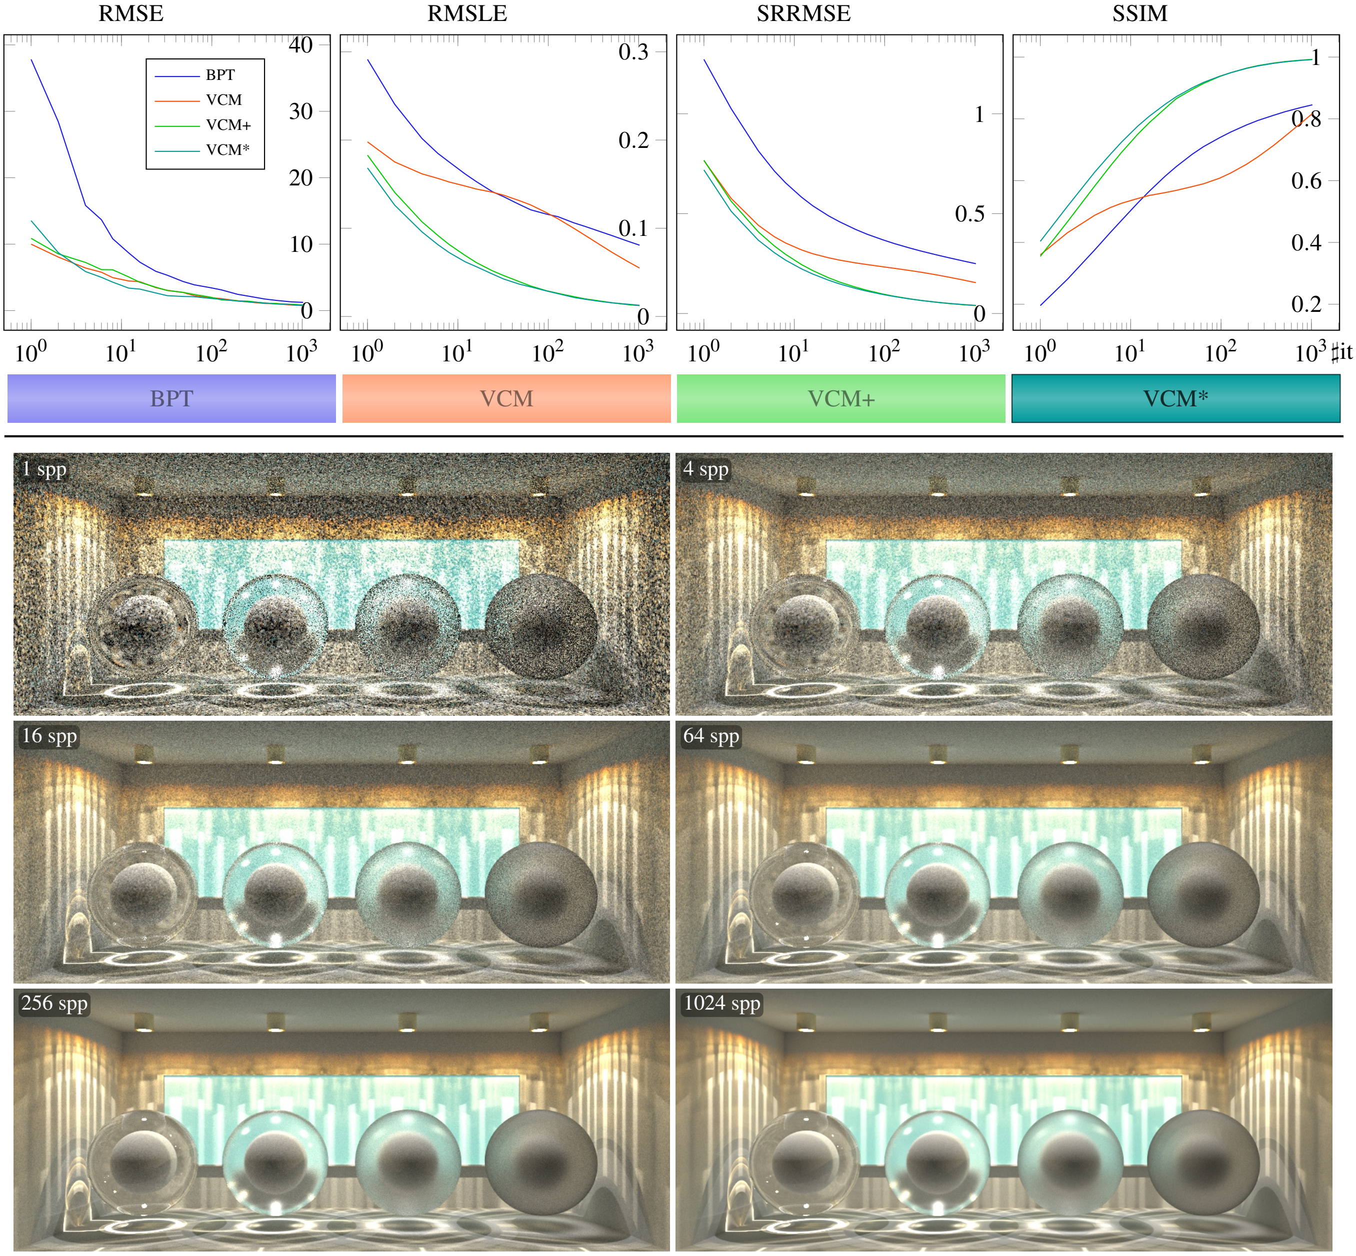

Figure 1: ***Spheres****: 1000x400 pixels, 4 point lights, 1 weak area light. BPT is not even close to convergence, the original VCM still contains salt noise but looks close to the final result wheres our method handles each light situation in the scene robustly. This is consistent with the error metrics in which our methods performs best in all cases.*

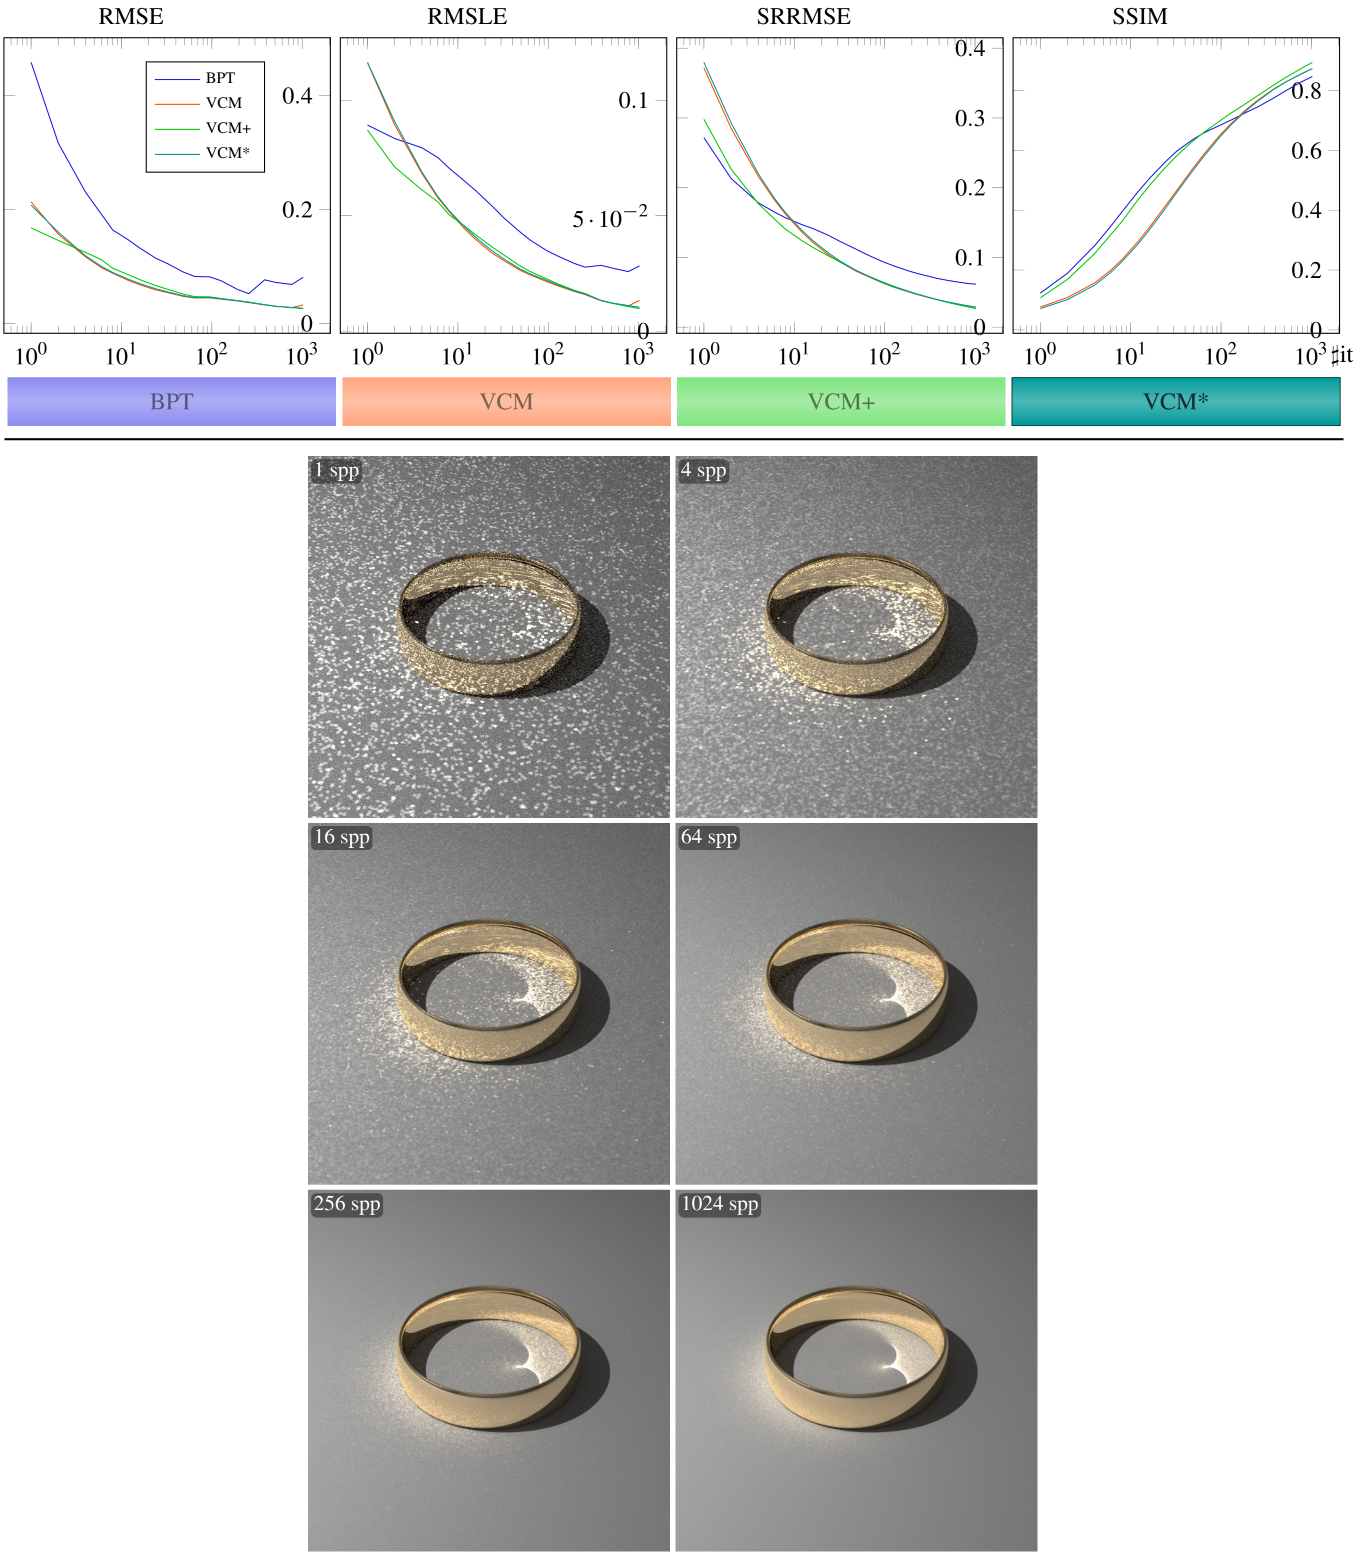

Figure 2: ***Ring****: 500x500 pixels, 1 point light. BPT misses reflected caustics. In this scene VCM+ improves the convergence on the diffuse surface and reduces the bias in the caustics compared to standard VCM and VCM*. VCM* performs the same as VCM.*

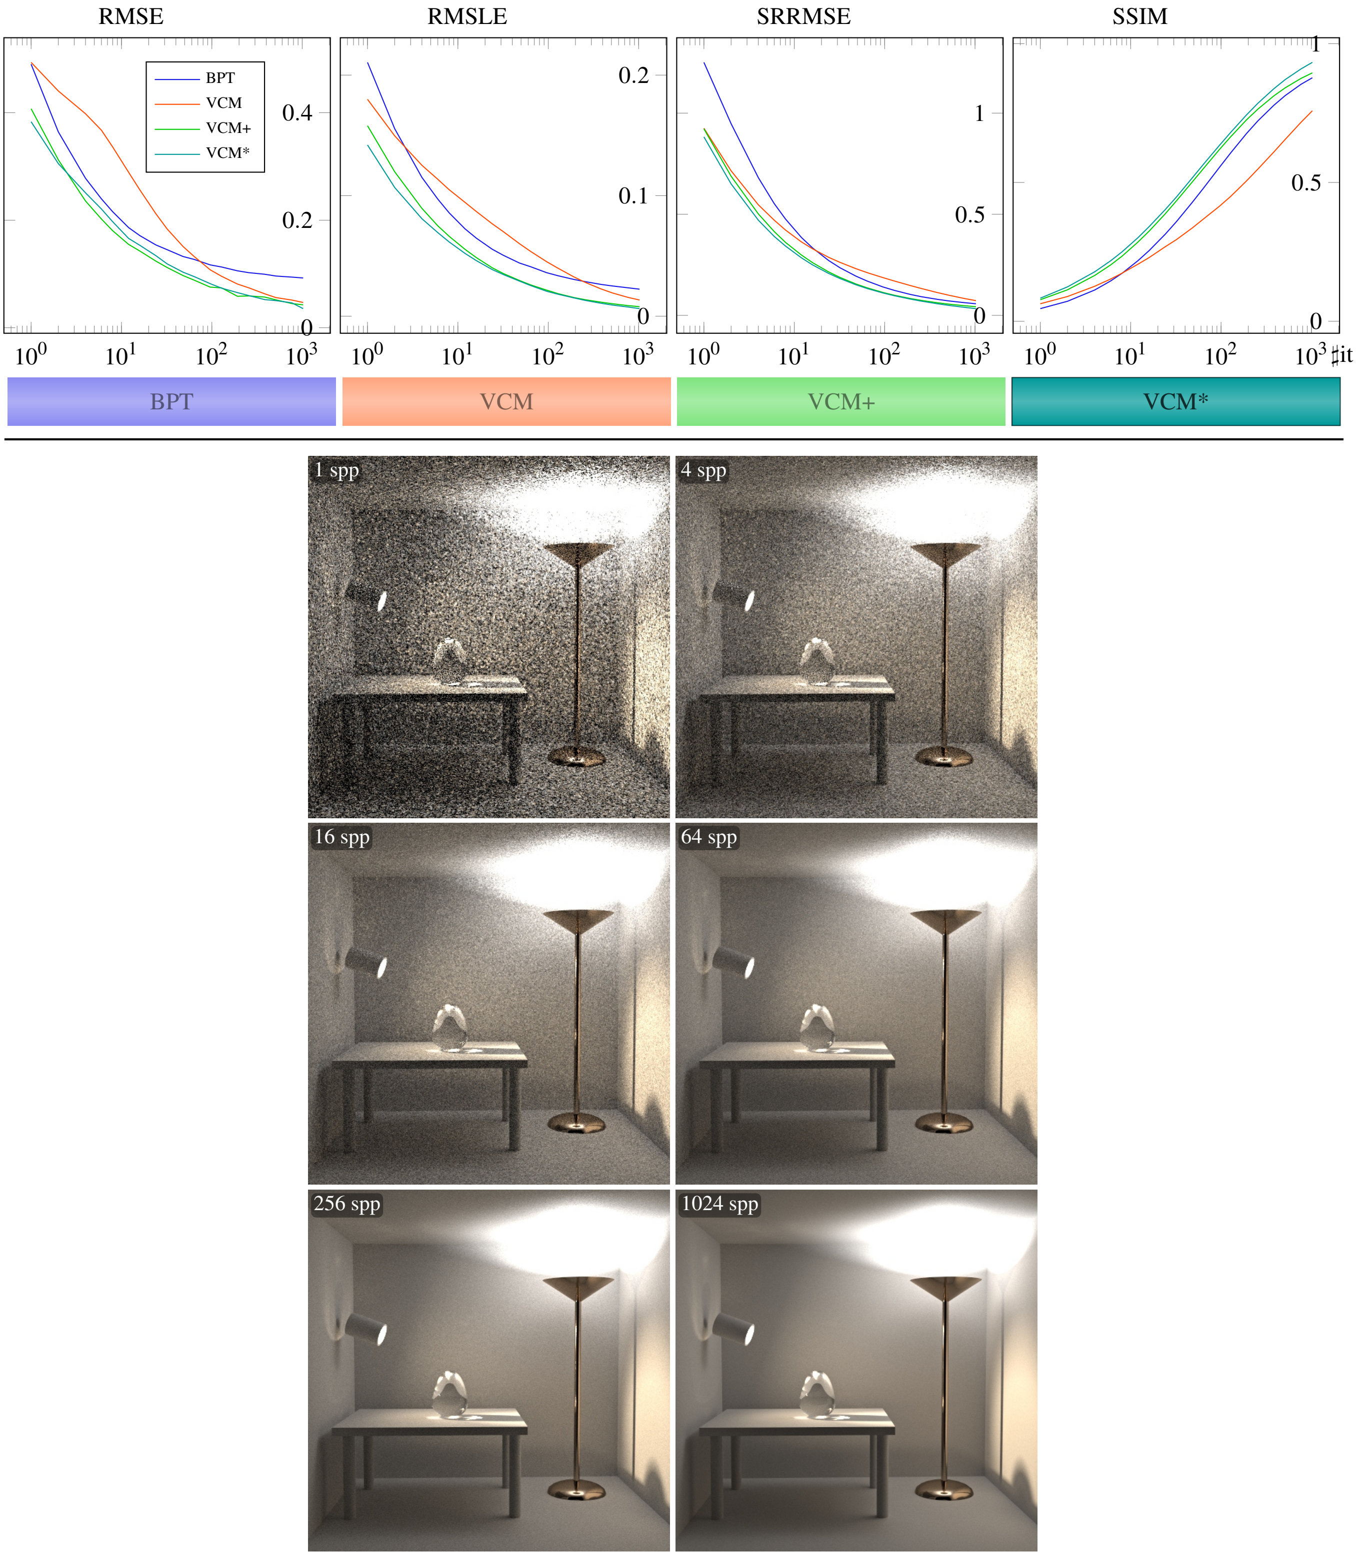

Figure 3: ***Veach-Bidir*** *from PBRT-v3 [PJH17]: 512x512 pixels, 2 areal lights. The cylinder on the left has again the problem of hard too reach geometry close to the light source leading to more noise in standard VCM than in our variants.*

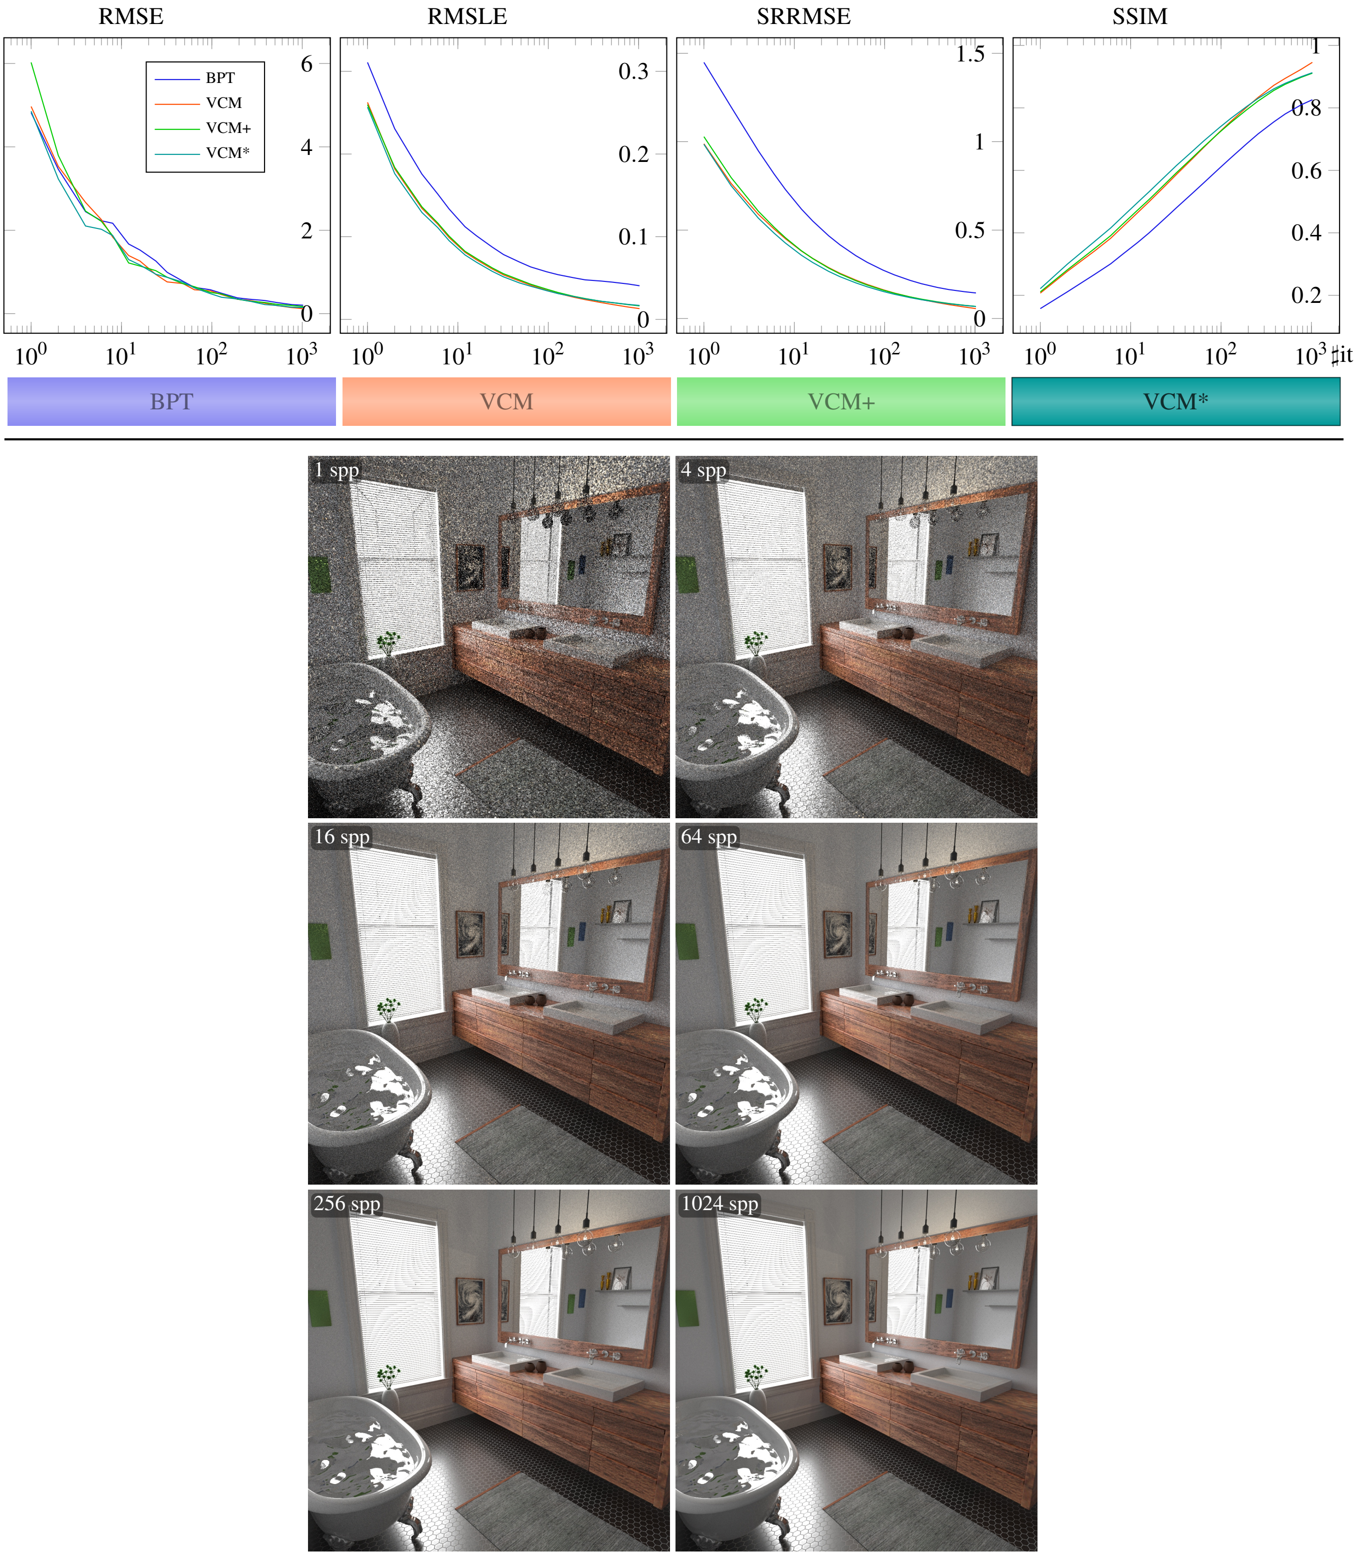

Figure 4: ***Contemporary Bathroom*** *from PBRT-v3 [PJH17], (Courtesy of Mareck): 1024x1024 pixels. All VCM variants are very similar and better than BPT. VCM* has the smallest error for most of the time. The reason why VCM is better than VCM* at the end is that VCM* choses merges on the infacing roller blind sides instead on outfacing ones. Due to light-bleeding within the gathering radius these have a higher bias.*

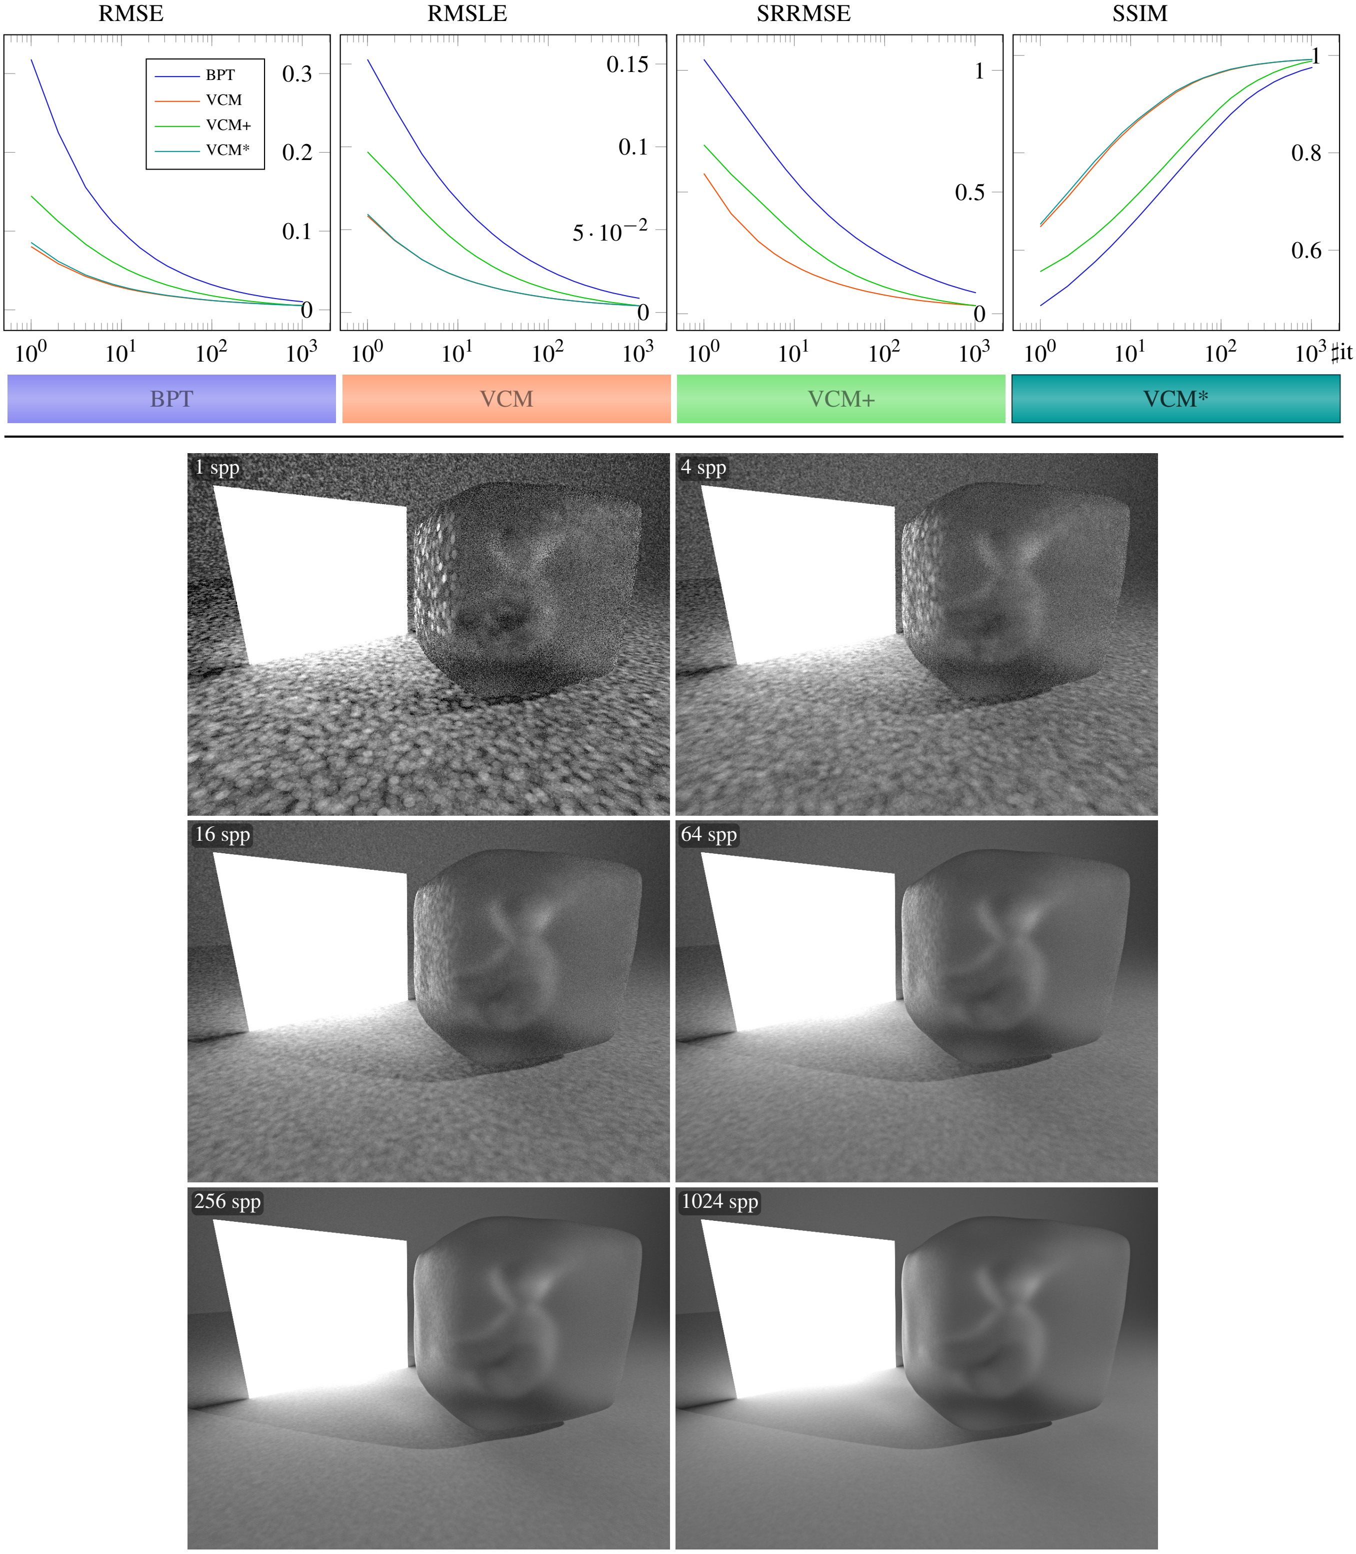

Figure 5: ***Bunnyduck***: *1024x768 pixels. VCM* is slightly worse than VCM, because the density of photons is overestimated. See the early iterations on the glass cube front face: (refractive) merges on this face are penalized due to high density from direct lighting of the second light source. VCM+ has higher frequency noise and looks subjectively better.*

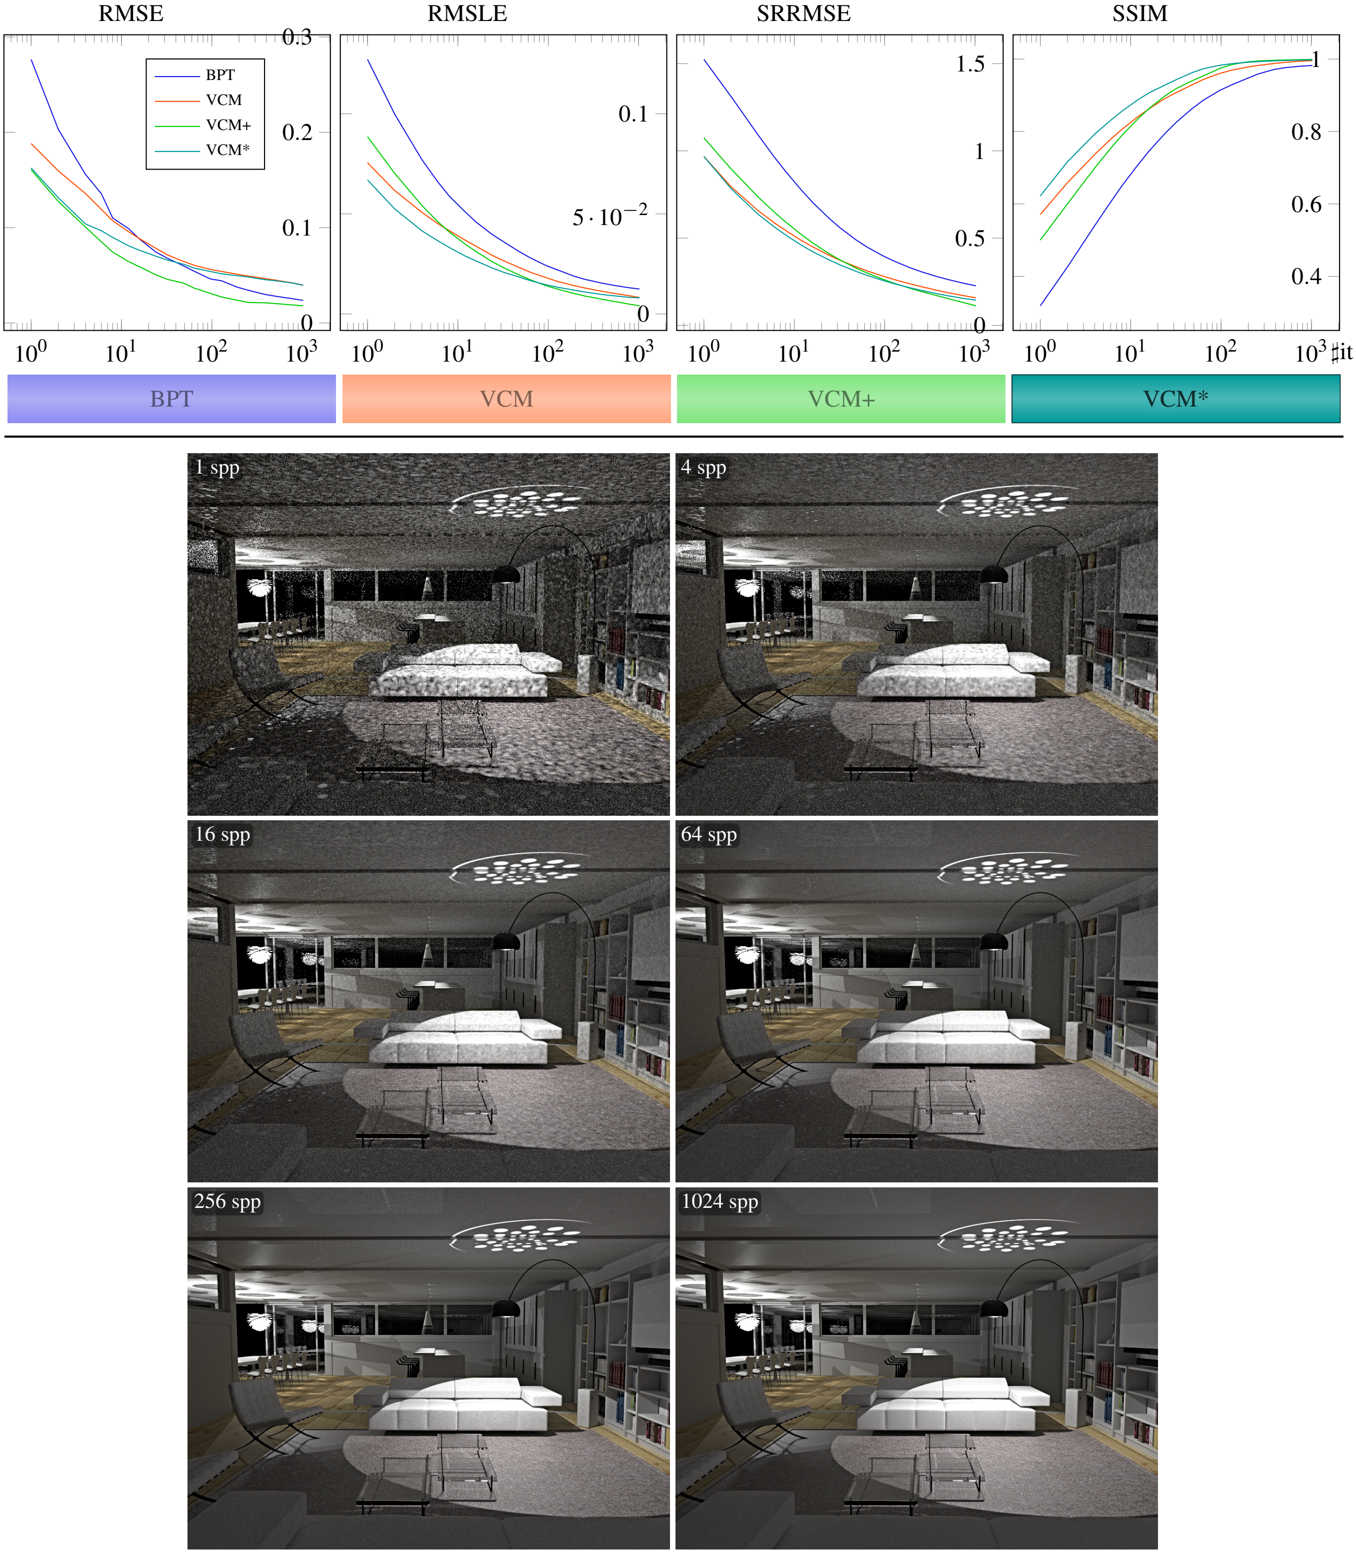

Figure 6: ***Villa*** *from PBRT-v3 [PJH17]: 1024x768 pixels. Robust in all methods. However, VCM shows some more fireflies (e.g. see center, 256spp, front left area; or the ceiling in less than 256spp images).*

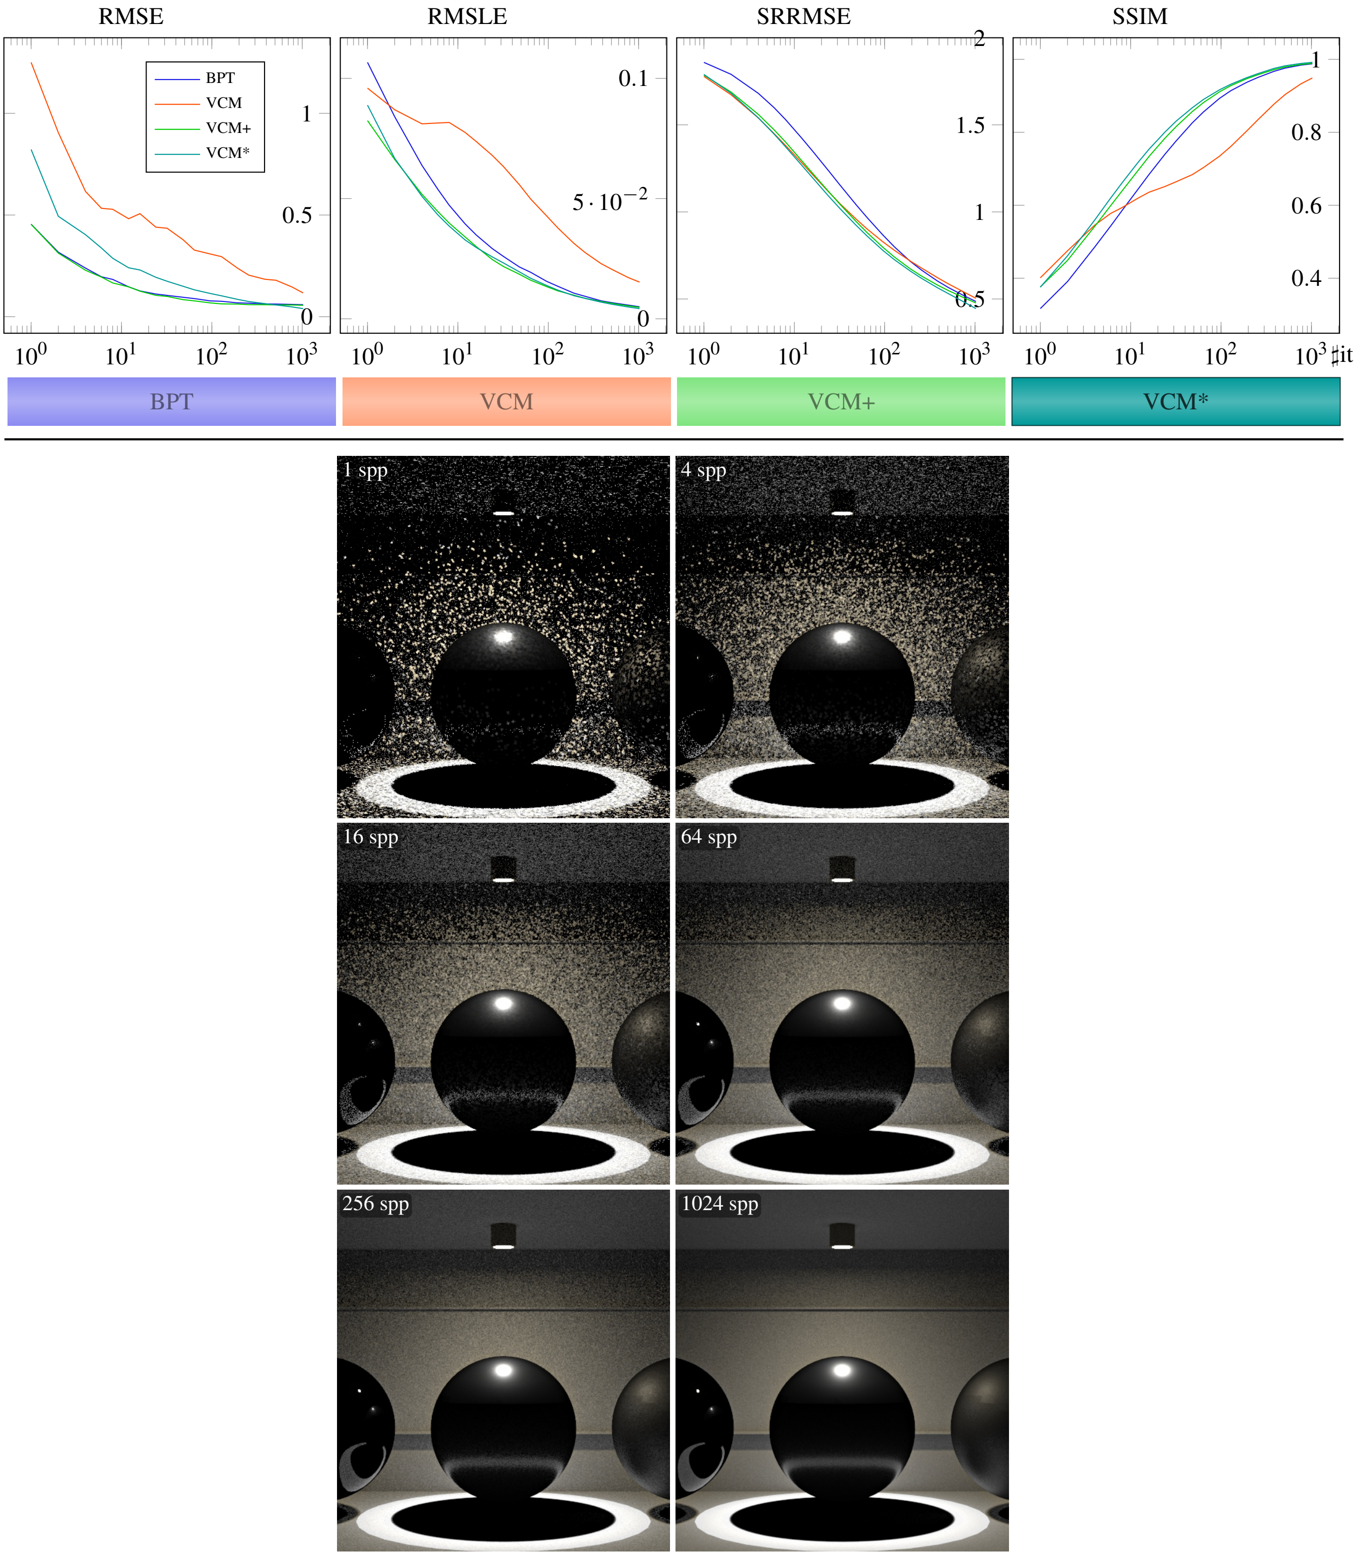

Figure 7: ***Spheres*** *with diffuse cylinder and max. path length of 3: 460x500 pixels. BPT and VCM+ win in RMSE, because the error is dominated by bias. VCM+ and VCM* clearly show their improved handling of the double-merge case over standard VCM. On the long run VCM* gets the best score in any measurement.*

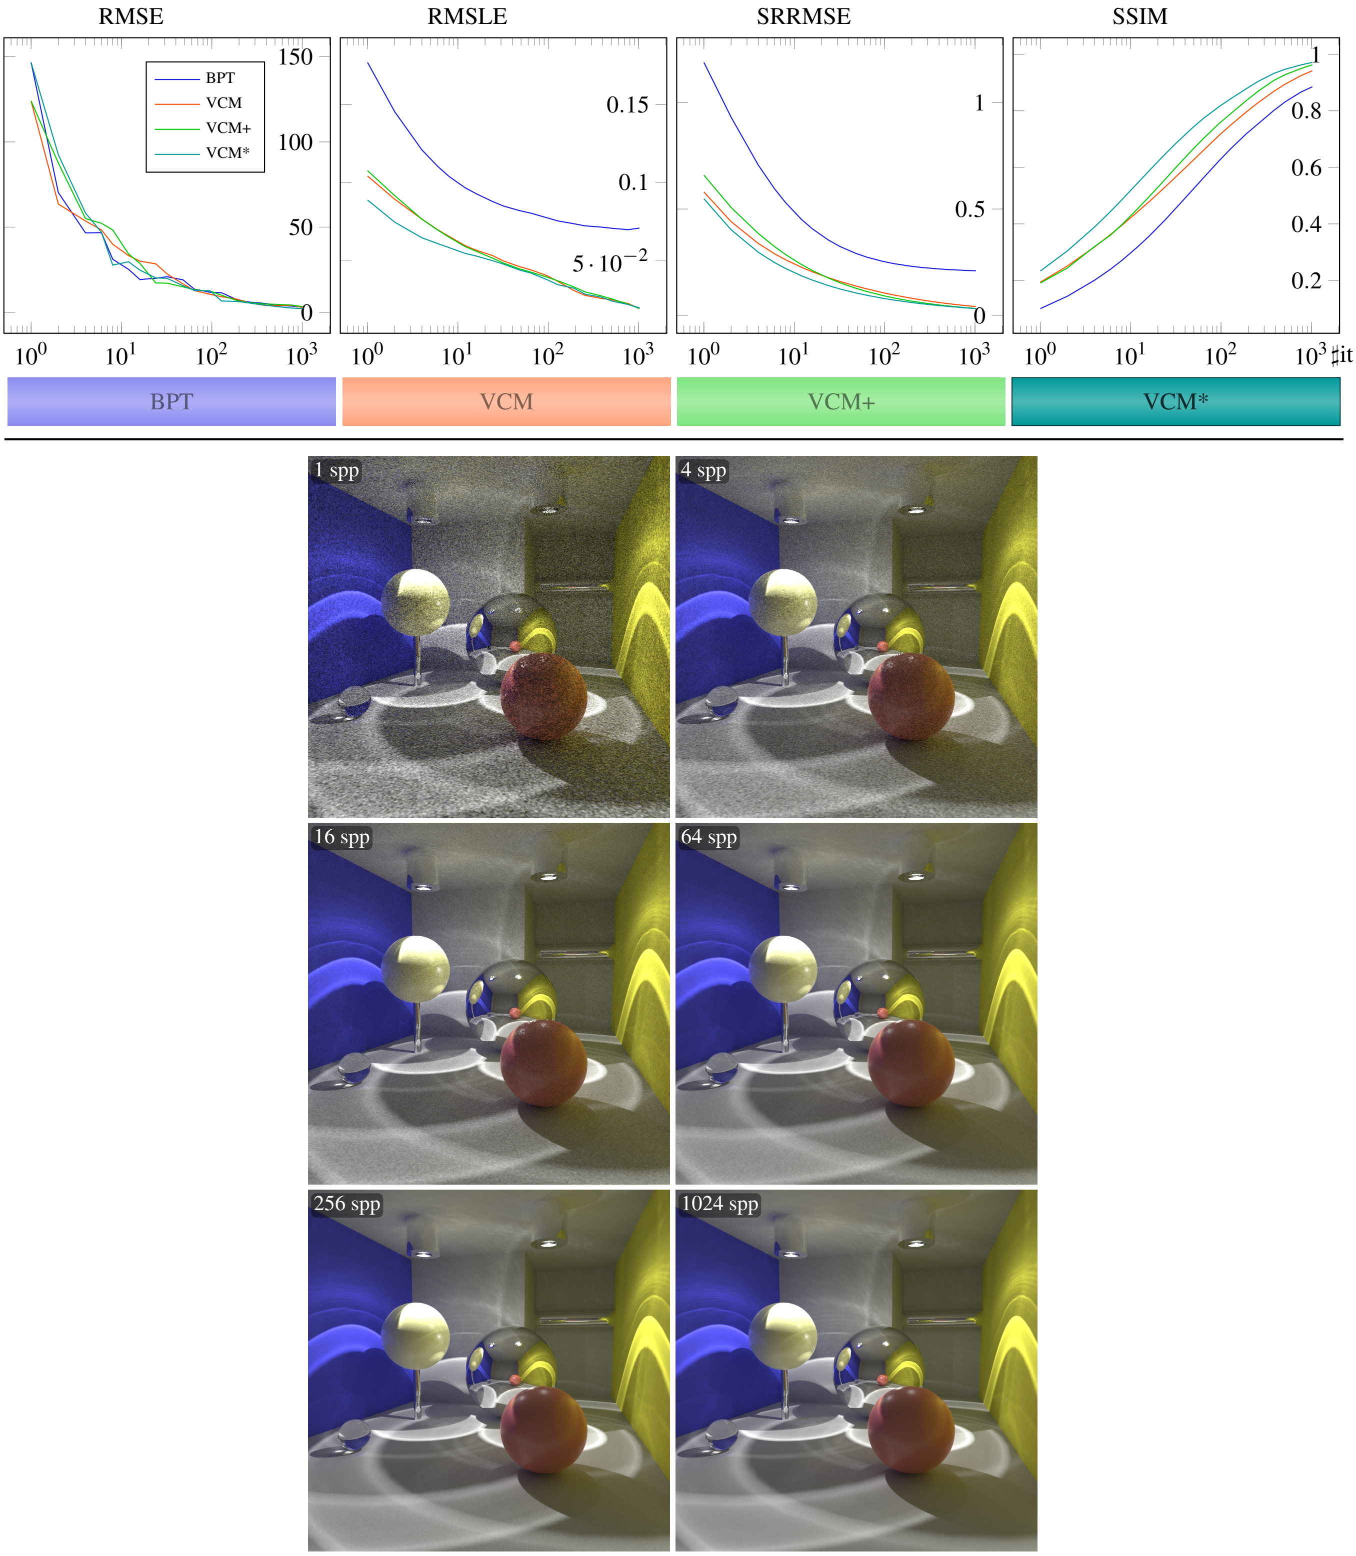

Figure 8: ***Mirrorballs*** *Courtesy of Toshya Hachisuka (taken from [Vev18]): 1024x1024 pixels. VCM and VCM* perform almost the same whereas BPT misses reflected caustics again. The small difference in SSIM is again due to higher noise in VCM (see dark floor areas).*

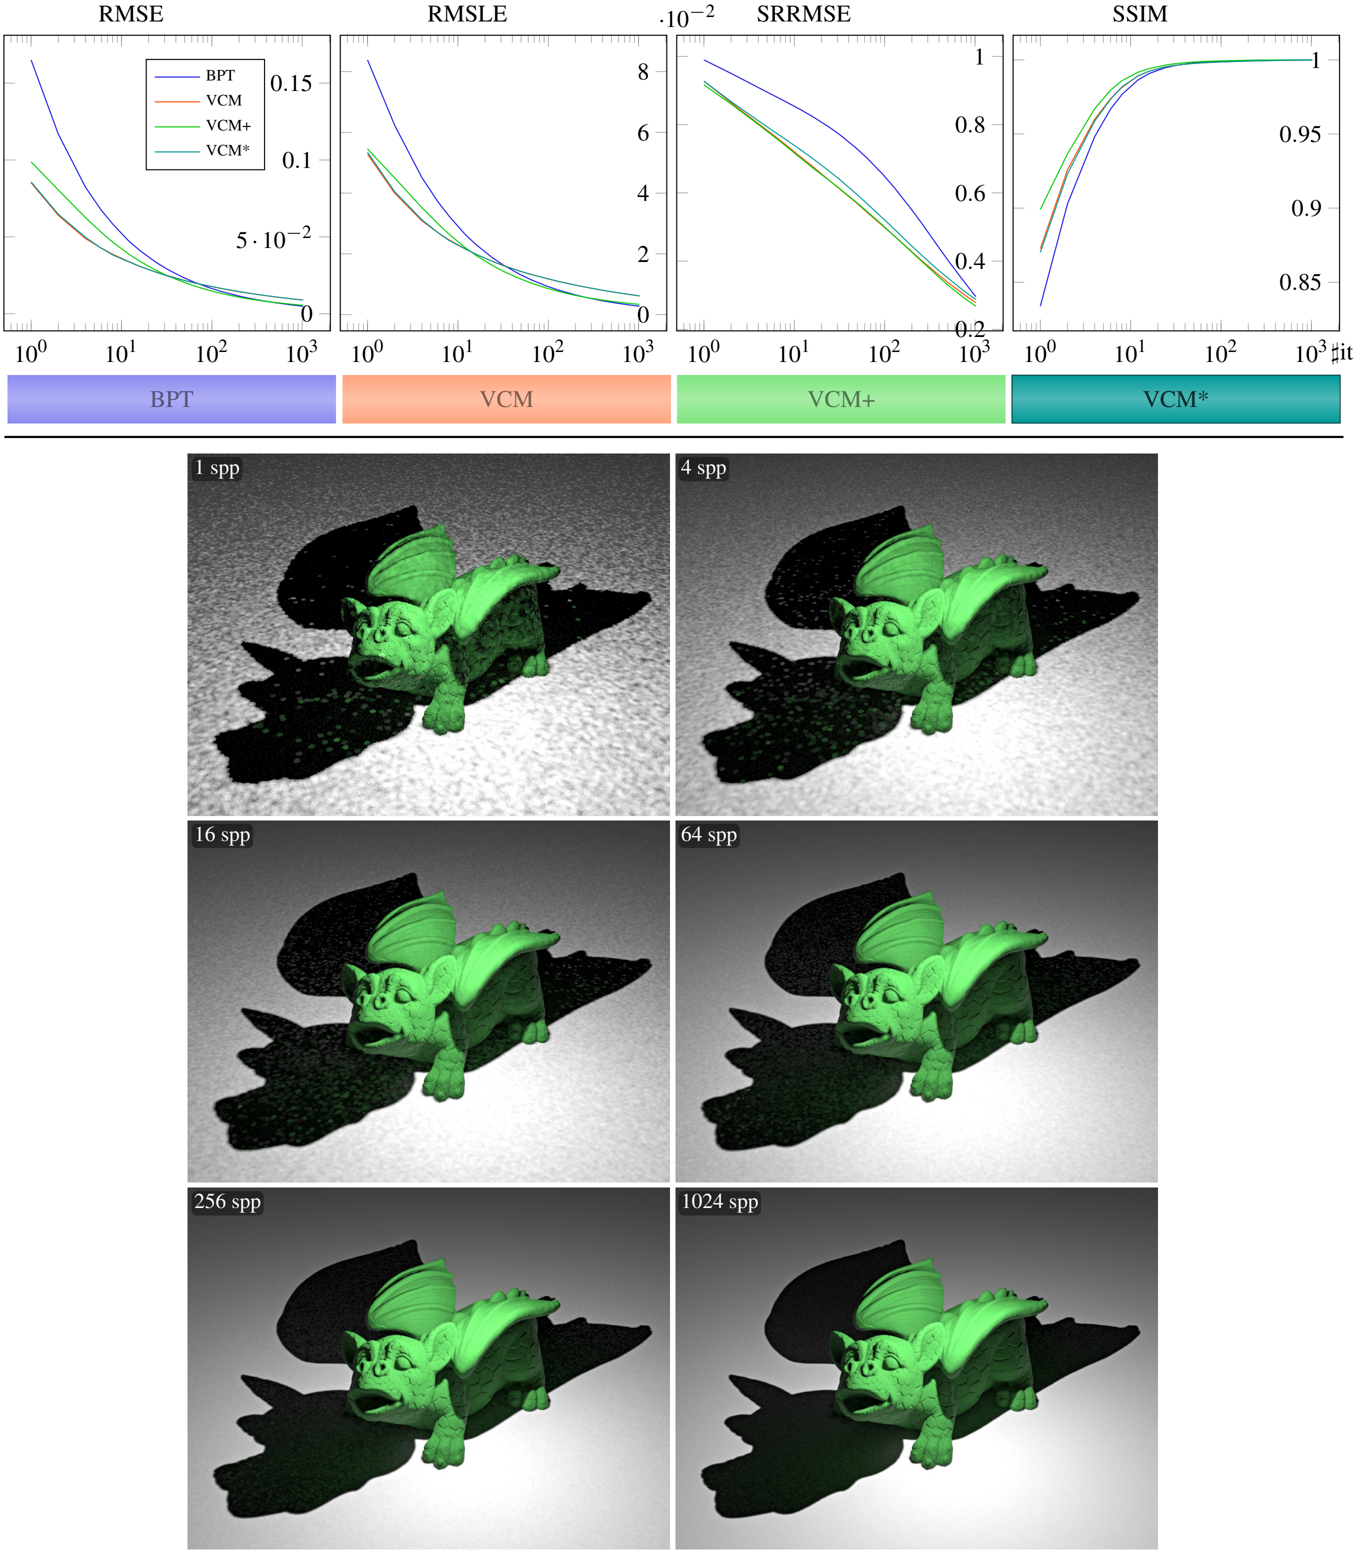

Figure 9: ***Dragon*** *from PBRT-v3 [PJH17] (Courtesy of Christian Schüller): 800x600 pixels. VCM+ produces less bias than standard VCM/VCM* (see shadow border). VCM/VCM* behave the same in this scene.*