Progressive Sampling Strategies for Disk Light Sources

Per Christensen
Pixar Technical Memo #18-02 — June 2018
Pixar Animation Studios

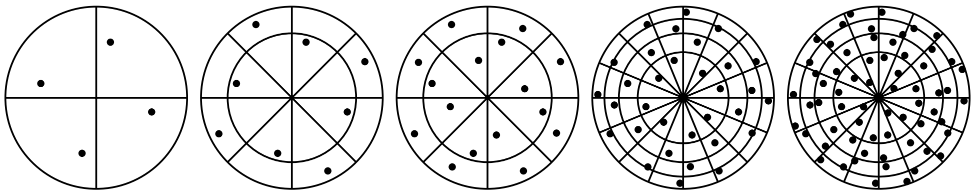

Figure 1: *The first 4, 8, 16, 32, and 64 samples from a progressive multi-jittered (0,2) sample sequence with the polar4 sampling strategy.*

Abstract

This technical memo compares six different strategies for sampling a disk area light source. We test padding the domain with zero values, rejection sampling, polar mapping, concentric mapping, and two new strategies we call polar4 and concentric4 mapping. We test the six strategies with nine different sample sequences on pixels with full illumination and partial shadow (penumbra), and find the best combination of sampling strategy and sample sequence. One of the new sampling strategies gives more than 3 times error reduction in some cases.

1 Introduction

Efficient sampling of area lights is a common problem in Monte Carlo and quasi Monte Carlo rendering. Disk lights are particularly interesting because sample sequences usually are generated on the unit square, so the samples have to be mapped to the disk or other sample strategies must be applied.

In this technical memo we seek a disk area sampling strategy with good stratification, giving low initial error and fast convergence. The goal is to approach the same convergence rates as seen for area sampling of square light sources: roughly $O(N^{-0.75})$ and $O(N^{-1.5})$ for penumbra regions and fully illuminated regions, respectively [Christensen et al. 2018]. It turns out that neither the standard polar mapping nor concentric mapping give the desired convergence and initial error properties, but a simple variation of polar mapping does.

We believe these strategies could also be beneficial for angular sampling of disk lights (which is equivalent to sampling of spherical ellipses) if combined with the transformations of Guillén et al. [2017].

2 Sample sequences

In this memo we test various disk light sampling strategies using the following sample sequences defined on the unit square: uniform random, best candidates [Mitchell 1991], Ahmed ART [Ahmed et al. 2017], Halton [1964], Sobol' (0,2) [1967], and progressive multi-jittered (0,2) [Christensen et al. 2018]. For brevity we call the last sequence *pmj02*. The Halton sequence is randomized with either rotations [Cranley and Patterson 1976] or random digit scrambling; the Sobol' sequence is randomized with rotations, xor-scrambling [Kollig and Keller 2002], or Owen scrambling [Owen 1997; Owen 2003]. Figure 2 shows examples of the first 500 samples from each of these sequences.

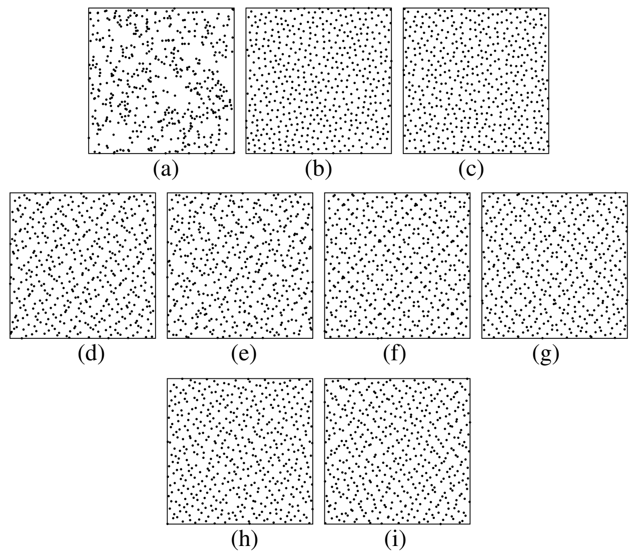

Figure 2: *First 500 samples from nine sample sequences: (a) uniform random; (b) best candidates; (c) Ahmed ART; (d) rotated Halton; (e) scrambled Halton; (f) rotated Sobol' (0,2); (g) xor-scrambled Sobol' (0,2); (h) Owen-scrambled Sobol' (0,2); (i) progressive multi-jittered (0,2).*

3 Sampling strategies

We test six different sampling strategies illustrated in Figure 3.

Pad-zero. With this strategy we scale the samples from the unit square to the square $[-1, 1]^2$ and pad the sample domain with zero values, that is, samples that fall outside the unit circle are assigned the sample value zero (corresponding to shadow, but without actually tracing a shadow ray). In Figure 3 (top left), three of the 16 sample values are zero. For this sampling strategy to converge to the correct value, the sample values must be multiplied by the ratio of the sample domain areas: $4/\pi$. Padding with zero may seem promising since it avoids any distortion of the sample strata, but unfortunately it introduces a discontinuity even if the illumination from the entire disk is smooth. As demonstrated in Christensen et al. [2018] and elsewhere, discontinuities reduce convergence rate to $O(N^{-0.75})$ so this strategy is not good.

Rejection sampling. Rejection sampling is a notoriously poor sampling strategy that is sometimes used as a last resort when better

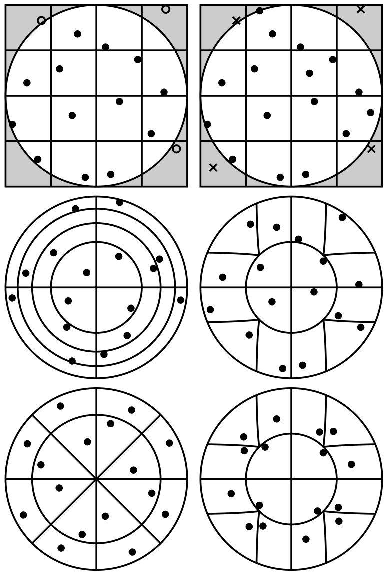

Figure 3: *Disk sampling strategies: pad-zero, rejection, polar, concentric, polar4, and concentric4.*

strategies are unavailable. We again scale the sample points from the unit square to $[-1,1]^2$, but if a sample falls outside the unit disk we pick the next sample (and repeat until the sample is inside the disk). In Figure 3 (top right) we have to take 20 samples to get 16 valid samples on the disk – four samples are rejected.

Polar map. The most common way to transform points (u, v) on the unit square to the unit sphere is by an area-preserving mapping to polar coordinates: $r = \sqrt{u}, \phi = 2\pi v$. This maps the entire unit square to the entire unit disk, so no samples are wasted and no discontinuities introduced. But unfortunately the strata have rather irregular shapes with poor aspect ratios (as can be seen in Figure 3 (middle left)).

Concentric map. To overcome the high distortion of polar mapping, Shirley and Chiu [1997] introduced concentric mapping. This mapping essentially "squeezes" the corners of the square into a disk shape, as shown in Figure 3 (middle right). Here square strata are mapped to strata with much more even aspect ratios. Intuitively this would seem like a much better mapping than the polar map, but as we will see, that is not the case for sampling of disk lights.

Polar4 map. This is a slight variation on polar mapping, shown in Figure 3 (lower left). Samples are mapped from the unit square only to the quarter-circle "wedge" in the first quadrant. But the same sample point is then rotated by 90, 180, and 270 degrees. As we will see, this reduces sampling error in a desirable way.

Concentric4 map. A slight variation on concentric mapping, similar to the polar4 map: samples are mapped from the unit square only to the quarter-circle "wedge" in the first quadrant and the same sample point is then rotated by 90, 180, and 270 degrees. This is shown in Figure 3 (lower right).

4 Tests

Figure 4 shows two teapots on a ground plane, illuminated by a disk area light source creating interesting penumbra regions with soft shadows. The disk light source is sampled with shadow rays shot from the surface point at the center of each pixel (no pixel jittering).

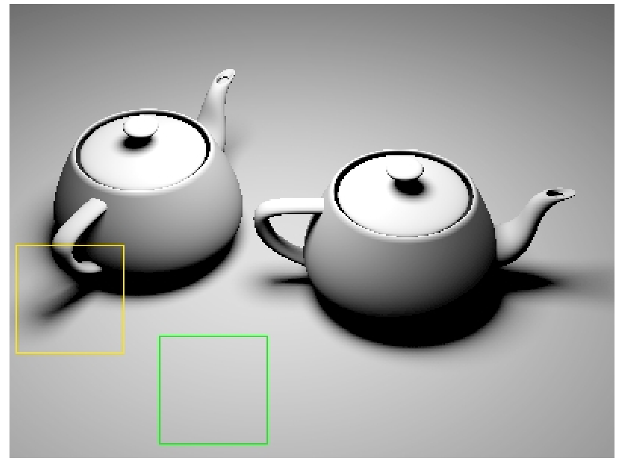

Figure 4: *Two teapots illuminated by a disk light source. The green square outlines a region with smooth illumination, the yellow square outlines a penumbra region with partial occlusion.*

The tests in this section are similar to the tests of square and rectangular area lights in Christensen et al. [2018], and are also similar in spirit to Ramamoorthi et al. [2012], but for sample sequences rather than sets. The plots in the following show rms error as function of the number of samples for combinations of six sampling strategies and nine sample sequences; each error curve is computed as the average of 100 sequences.

4.1 Fully illuminated region

Figure 5 shows error plots for pixels in the smooth, fully illuminated image region marked with a green square. In this region, the only variation in illumination is due to varying angle and distance to points on the light source, not occlusion.

As expected, the pad-zero and rejection sampling strategies have high error and poor convergence: even though the illumination in this region is smooth, the error convergence is only $O(N^{-0.75})$ for the best sample sequences.

For the polar mapping, random and best candidates converge at $O(N^{-0.5})$ as usual. Rotated and scrambled Halton, and rotated and xor-scrambled Sobol' (0,2), converge as roughly $O(N^{-1})$. Pmj02 and Owen-scrambled Sobol' (0,2) converge faster at roughly $O(N^{-1.25})$ at sample counts that are powers of two. This is consistent with results for square lights (except that the last rate is slower than for square light sources).

But at the same time, something interesting and unexpected is happening: *the error of xor-scrambled Sobol' starts out much lower than the error for the other sequences*, and it isn't until 2048 samples that the faster convergence rate of pmj02 and Owen-scrambled Sobol' (0,2) catches up. We believe this is because the polar mapping is antisymmetric – systematic shifts of the samples on the unit square will move half the samples on the disk in one direction and half in the opposite direction – and this works well with the xor scrambling. This result is similar to the observation of Ramamoorthi et al. [2012] that uniform jittering is a surprisingly good sample set for disk lights (although uniform jittering is a very poor sample set in many other settings). This also explains why Kensler [2013] found polar mapping to be slightly better than concentric mapping for his correlated multi-jittered sample sets.

For the concentric mapping, random and best candidates converge as $O(N^{-0.5})$, while most other sample sequences converge as

$O(N^{-1})$. The convergence rate for pmj02 and Owen-scrambled Sobol' (0,2) is around $O(N^{-1.15})$ at powers of two. Here there are no surprises: xor scrambling is clearly worse than Owen scrambling.

Given these insights, we propose a new sampling strategy for disk lights: use pmj02 or Owen-scrambled Sobol' (0,2) samples with polar mapping, but use each sample four times with the polar angle ϕ rotated 0, 90, 180, and 270 degrees. This new sampling strategy utilizes the antisymmetric nature of the polar mapping: if a sample point is, e.g., on the upper edge of its stratum and jittered away from the illuminated point, then the second rotated version of that sample point will be jittered toward the illuminated point. We call this strategy polar4.

As it can be seen in the error plot for polar4, this strategy results in a combination of the low initial error of polar mapped xor-scrambled Sobol' with the fast convergence rate of pmj02 and Owen-scrambled Sobol' (0,2). The convergence rate is actually better at roughly $O(N^{-1.4})$ at powers of two.

With this sampling strategy, even the error for uniform random samples (where we now use the same random sample four times, offset by 90 degrees) is quite a bit lower than when using other sampling strategies.

With 256 samples, the error is 8.5×10^{-5} with xor scrambled Sobol' and polar mapping, but only 2.5×10^{-5} with pmj02 (or Owen-scrambled Sobol' (0,2)) with the polar4 mapping, a reduction by 3.5 times. For another comparison, reaching an error below 10^{-4} requires 256 samples for the polar strategy, but only 128 samples for polar4; reaching an error below 10^{-5} requires 2048 samples for the polar strategy, but only 512 samples for polar4.

The polar and polar4 disk sampling strategies are reminiscent of hemisphere sampling for irradiance, where Ward and Shakespeare [1998] recommended having roughly π times as many strata in the polar direction than the radial direction for jittered sample sets. With polar4 we have 4 times as many strata in the polar direction.

We also tested the concentric4 disk sampling strategy. The sample sequences have (roughly) the same convergence rate as for concentric. The error for random and best candidates is lower, but for rotated Halton and rotated Sobol' (0,2) it is higher. Unfortunately the pmj02 and Owen-scrambled Sobol' (0,2) sequences do not benefit significantly from concentric4 instead of concentric mapping, their errors only change a few percent.

4.2 Penumbra region

Figure 6 shows error plots for pixels in the penumbra region marked with a yellow square in Figure 4. Here all sample sequences have the same convergence rate as in the penumbra region for a square light source: random and best candidates converge as $O(N^{-0.5})$ and all other sequences converge as $O(N^{-0.75})$.

The pad-zero and rejection sampling strategies have very high error and poor convergence, as expected. One interesting detail is that the rotated Sobol' sequence performs poorly with the pad-zero strategy.

Ahmed's ART sequence has higher error with polar mapping than with concentric, but for the other sequences, polar or concentric mapping does not seem to matter in the penumbra region.

The polar4 and concentric4 strategies reduce error for the uniform random, best candidates, and Ahmed ART sequences quite significantly. But for the pmj02 and Sobol' (0,2) sequences, the reduction is only a few percent.

5 Conclusion

We have analyzed the advantages and disadvantages of the standard approaches to sampling disk area lights, and have presented two simple variations – one of which has both low initial error and fast convergence at fully illuminated points and no disadvantages for points in penumbra regions.

The winning combination is the polar4 strategy with well-stratified samples – either stochastic pmj02 samples or quasi-random Owen-scrambled Sobol' (0,2) samples. These two combinations give a convergence rate – in fully illuminated areas – of $O(N^{-1.4})$ which is close to the $O(N^{-1.5})$ that can be obtained for square light sources. Using the polar4 mapping instead of the standard polar (or concentric) mapping gives up to a 3.5 times error reduction at fully illuminated points.

We have also found an explanation for some seemingly counter-intuitive results reported by other researchers [Ramamoorthi et al. 2012; Kensler 2013] testing various sampling strategies and sample sets for disk light sources.

Acknowledgements

This technical memo grew out of the supplemental material for the paper "progressive multi-jittered sample sequences" which was co-authored by Andrew Kensler and Charlie Kilpatrick.

References

AHMED, A., NIESE, T., HUANG, H., AND DEUSSEN, O. 2017. An adaptive point sampler on a regular lattice. *ACM Transactions on Graphics (Proc. SIGGRAPH) 36*, 4.

CHRISTENSEN, P., KENSLER, A., AND KILPATRICK, C. 2018. Progressive multi-jittered sample sequences. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering) 37*, 4.

CRANLEY, R., AND PATTERSON, T. 1976. Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis 13*, 6, 904–914.

GUILLÉN, I., UREÑA, C., KING, A., FAJARDO, M., GEORGIEV, I., LÓPEZ-MORENO, J., AND JARABO, A. 2017. Area-preserving parameterizations for spherical ellipses. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering) 36*, 4, 179–187.

HALTON, J. 1964. Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the ACM 7*, 12, 701–702.

KENSLER, A. 2013. Correlated multi-jittered sampling. Tech. Rep. 13-01, Pixar Animation Studios.

KOLLIG, T., AND KELLER, A. 2002. Efficient multidimensional sampling. *Computer Graphics Forum (Proc. Eurographics) 21*, 3, 557–563.

MITCHELL, D. 1991. Spectrally optimal sampling for distribution ray tracing. *Computer Graphics (Proc. SIGGRAPH) 25*, 4, 157–164.

OWEN, A. 1997. Monte Carlo variance of scrambled net quadrature. *SIAM Journal on Numerical Analysis 34*, 5, 1884–1910.

OWEN, A. 2003. Quasi-Monte Carlo sampling. In *SIGGRAPH Monte Carlo Ray Tracing Course Notes*. ACM.

RAMAMOORTHI, R., ANDERSON, J., MEYER, M., AND NOWROUZEZAHRAI, D. 2012. A theory of Monte Carlo visibility sampling. *ACM Transactions on Graphics 31*, 5.

SHIRLEY, P., AND CHIU, K. 1997. A low distortion map between disk and square. *Journal of Graphics Tools 2*, 3, 45–52.

SOBOL', I. 1967. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics 7*, 4, 86–112.

WARD LARSON, G., AND SHAKESPEARE, R. 1998. *Rendering with Radiance: The Art and Science of Lighting Visualization*. Morgan Kaufmann.

A Pseudocode

Here we list pseudocode for the polar4 mapping. It is very simple. The input is an (unsigned) integer sample number s (which indexes into a table of two-dimensional sample points in the unit square $[0, 1]^2$). Output is a sample position on the unit disk (in polar coordinates (r, ϕ)).

procedure POLAR4(s, r, ϕ)
    // Reuse each sample point four times
    u $\leftarrow$ sample[s div 4].x
    v $\leftarrow$ sample[s div 4].y
    // Map point from unit square to polar coords on quarter-disk
    r $\leftarrow \sqrt{\text{u}}$
    $\phi \leftarrow$ 0.5 * π * v   // $\phi \in [0, \pi/2]$
    // Rotate sample by 0, π/2, π, or 3/2π
    $\phi \leftarrow \phi$ + 0.5 * π * (s mod 4)
end procedure

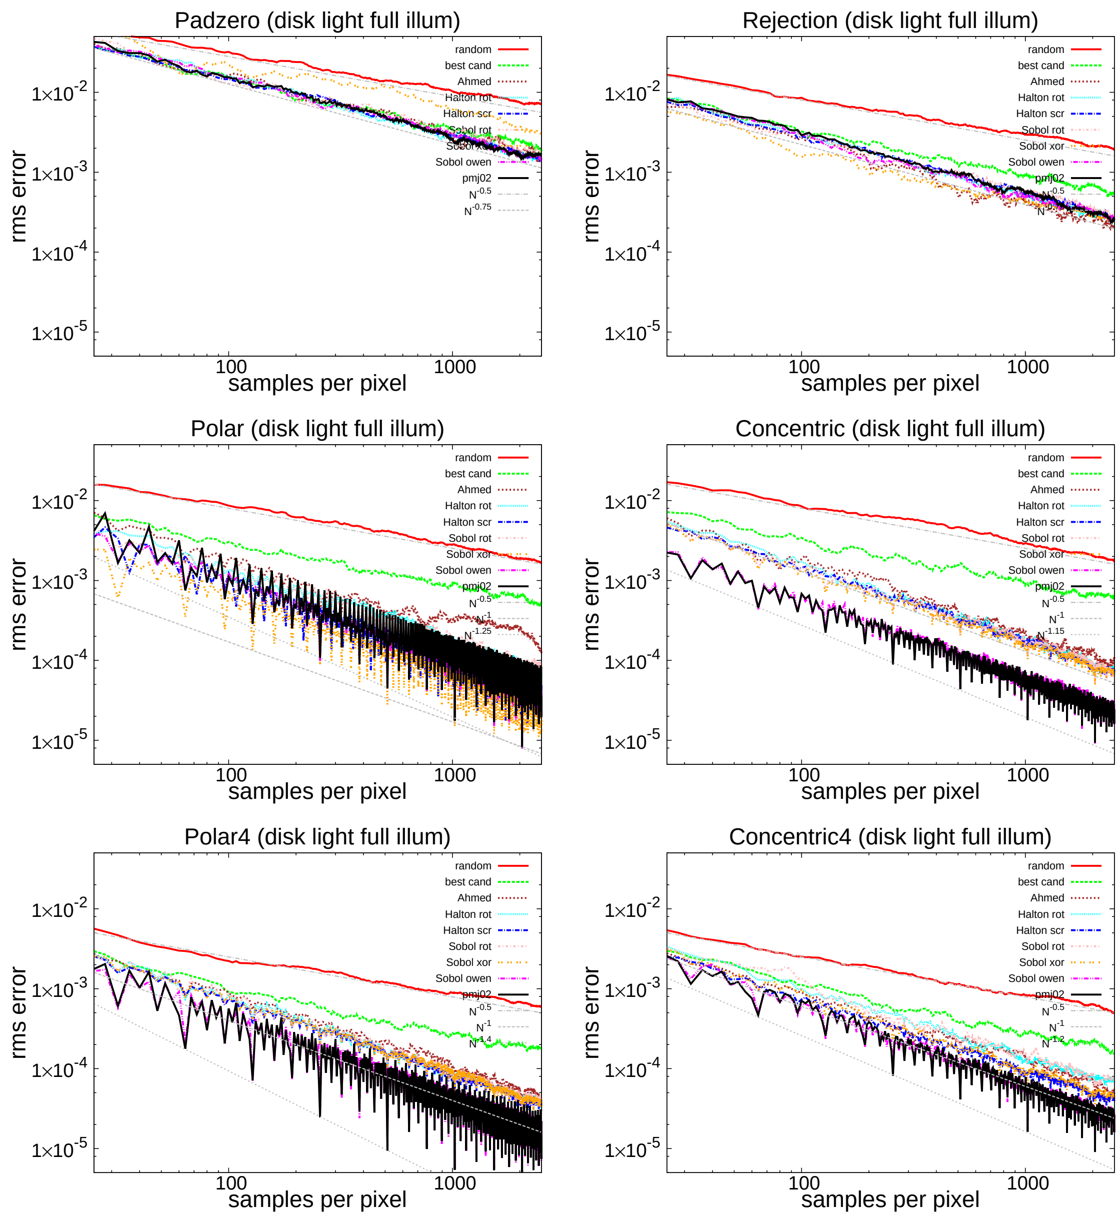

Figure 5: *rms error for 25–2500 samples per pixel of a disk light source, fully illuminated region. Disk sampling strategies: pad-zero, rejection, polar, concentric, polar4 and concentric4.*

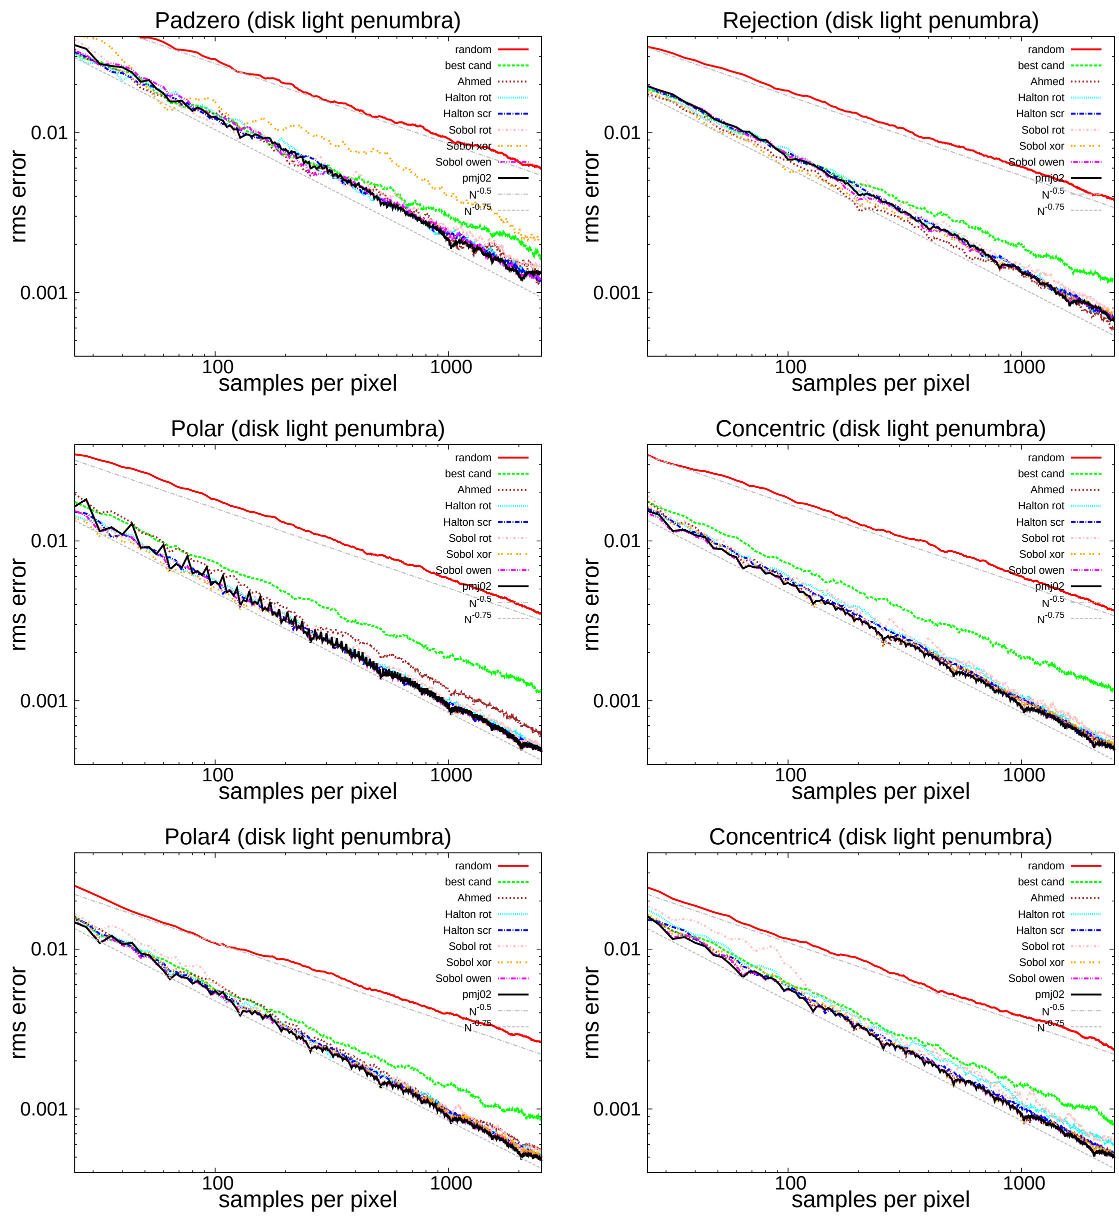

Figure 6: *rms error for 25–2500 samples per pixel of a disk light source, penumbra region. Pad-zero, rejection, polar, concentric, polar4, and concentric4 sampling strategies.*