Efficient Spatial Resampling Using the PDF Similarity

Yusuke Tokuyoshi

ARR
Advanced Rendering Research Group

AMD

Problem on Spatiotemporal Reservoir Resampling (ReSTIR) [Bitterli et al. 2020]

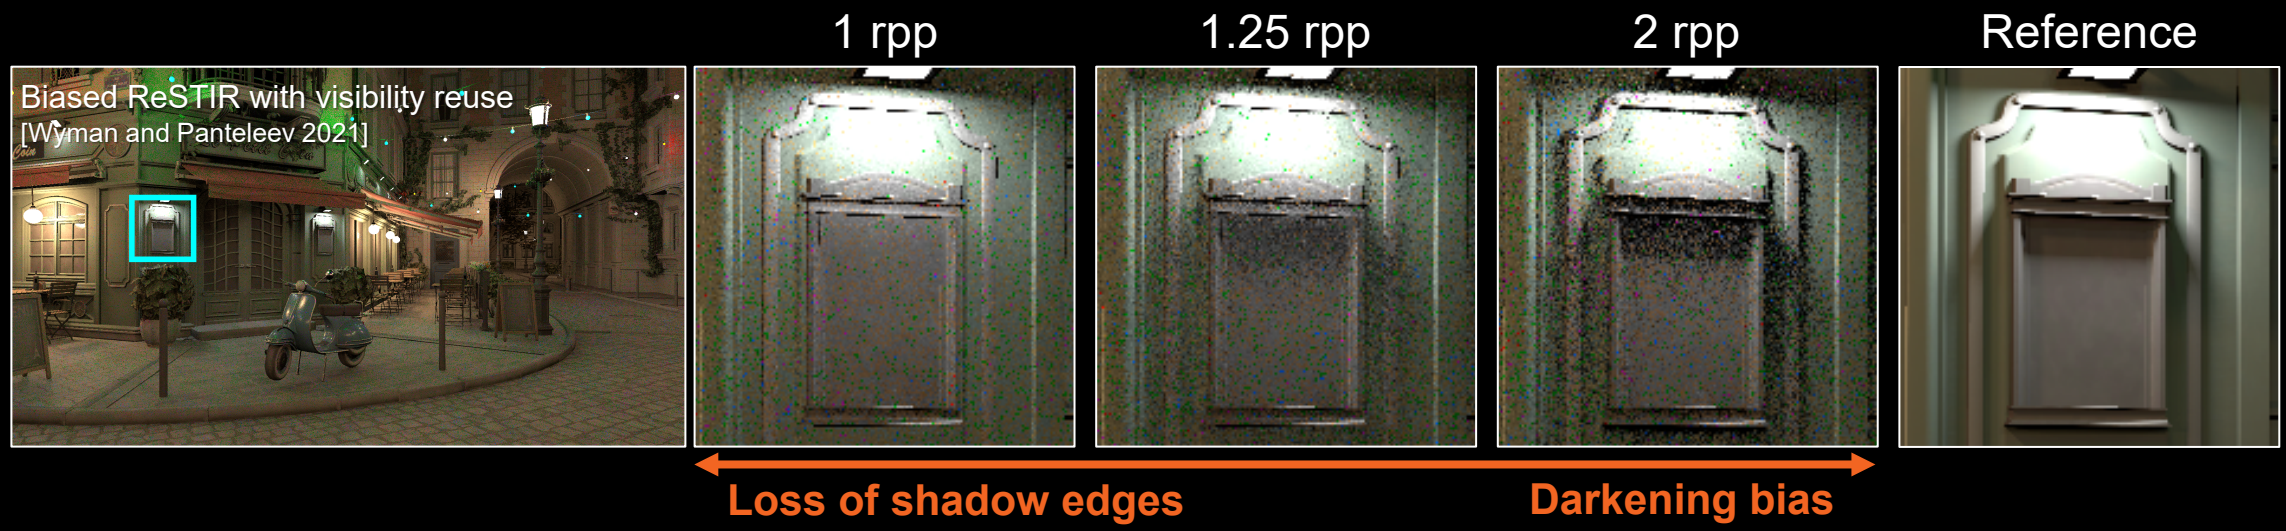

- Reusing samples is not always efficient for highly detailed scenes
  - Geometry edges, normal maps, spatially varying materials, and shadow edges
- Variance and **bias** due to the difference of target PDFs between pixels
  - Visibility reuse using 2 shadow rays per pixel (rpp) introduces a darkening bias ☹

AMD

Reject samples that have dissimilar PDFs

Previous Rejection Heuristics

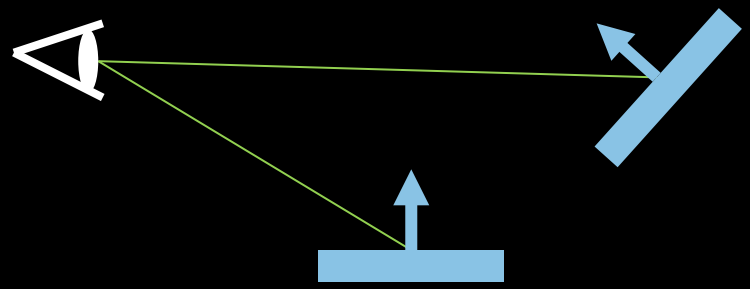

Similarity in geometry (depth & surface normal) between pixels

- Use a similarity of geometry [Bitterli et al. 2020]
- Use a roughness parameter and edge length [Lin et al. 2022]
- Ignore the visibility term in the PDF ☹
- Does not support arbitrary materials ☹

AMD

Our Rejection Heuristic

Previous (2 rpp)

Ours (1.25 rpp)


Reference

- Similarity of the PDF shapes including the visibility and BSDF terms
- Reduce the bias and variance around shadow edges and material boundaries
- Can reduce the number of shadow rays while preserving shadow edges

Our Contributions

- Rejection heuristic based on the similarity of PDF shapes
- Efficient PDF shape approximation using a von Mises-Fisher (vMF) distribution for single-bounce path connections (e.g., direct illumination)
- Temporal estimation for the vMF approximation
- Hybrid method with existing heuristics for animation
- Demonstration for ReSTIR variants with different visibility computation methods

AMD


CPU 0.123 ms, GPU 11.950 ms, 60 Hz
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Previous (2 rpp) ReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
2 rpp
Restaurant
Le Petit Coin
Previous rejection method


CPU 0.132 ms, GPU 12.196 ms, 60 Hz
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Ours (2 rpp) ReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
2 rpp
Restaurant
Le Petit Coin
Our rejection method


CPU 0.095 ms, GPU 11.010 ms,
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Previous (1 rpp) ReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
1 rpp
Restaurant
Previous method loses hard contact shadows due to spatial reuse of visibilities


CPU 0.086 ms, GPU 11.191 ms,
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Ours (1 rpp) ReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
1 rpp
Restaurant
Using our rejection method, we can render hard contact shadows without tracing two rays

Our Rejection Heuristic

Normalized Target PDFs

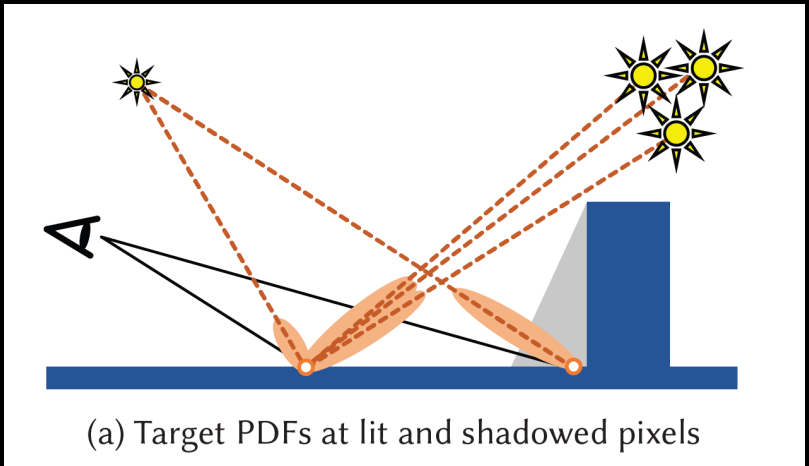

(a) Target PDFs at lit and shadowed pixels

- Use normalized target PDF shapes instead of the unnormalized target distributions
  - Dominant light directions are often different between lit and shadowed pixels
- Expensive ☹
  - Cannot obtain the normalization factor analytically
  - Infeasible to compare the exact PDF shapes between pixels

AMD

Approximation for the Single-Bounce Case

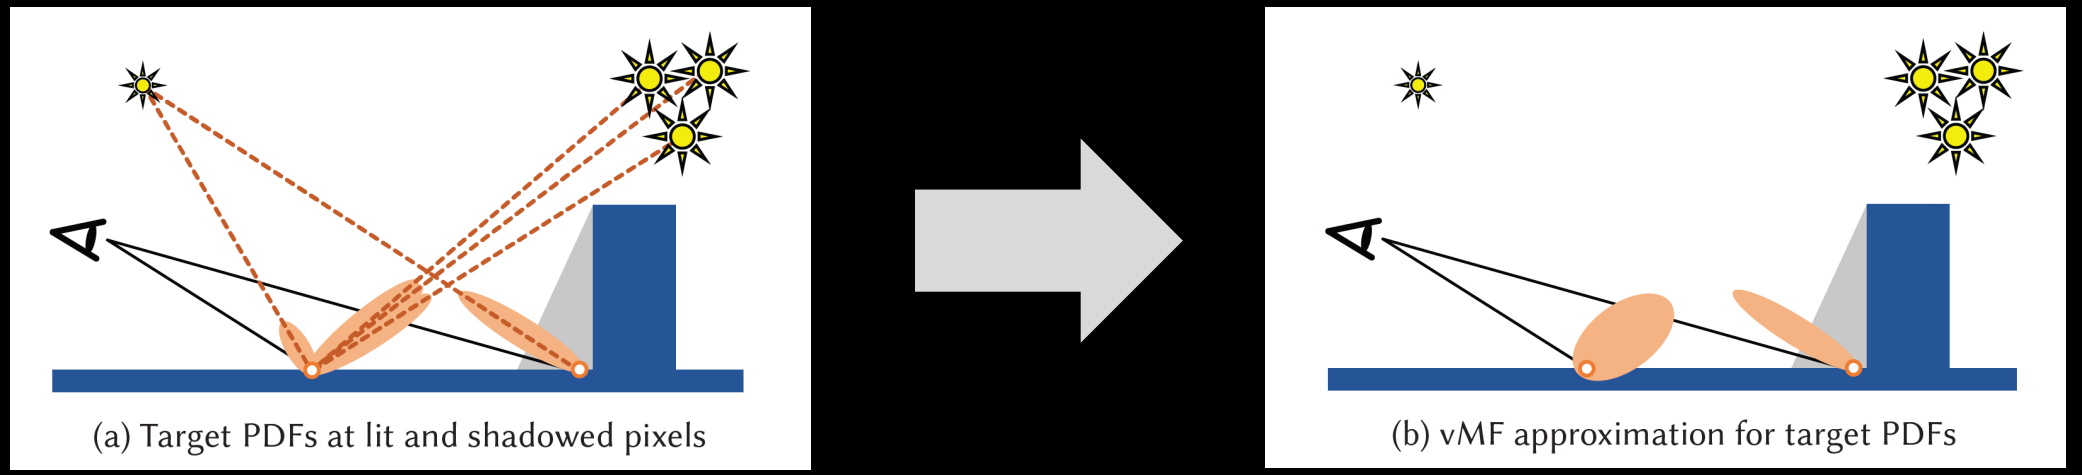

- Single-bounce PDF can be expressed with a spherical PDF
- Approximate this spherical PDF with a von Mises-Fisher (vMF) distribution [1953]
  - A.k.a. normalized spherical Gaussian [Tsai and Shih 2006; Wang et al. 2009]
  - Isotropic single-lobe distribution represented with axis and sharpness parameters

AMD

vMF Sharpness

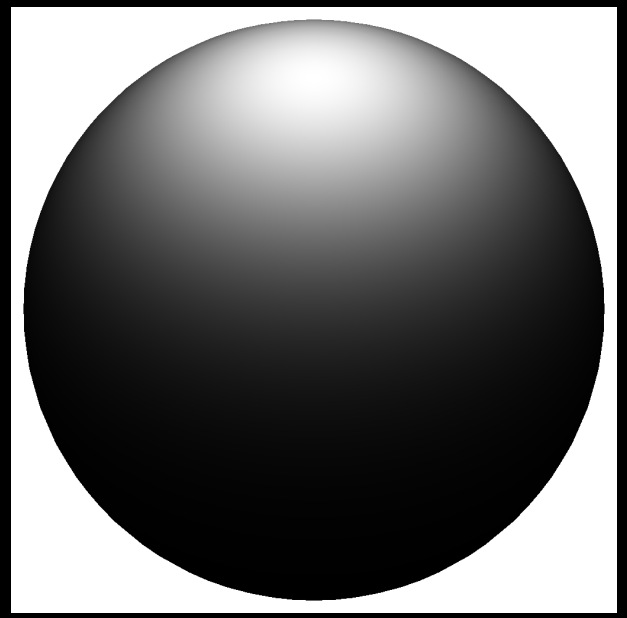

Sharpness=10

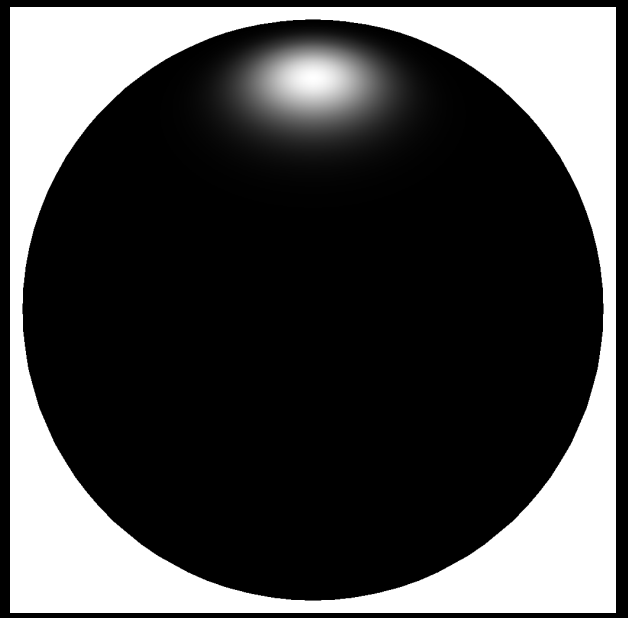

Sharpness=100

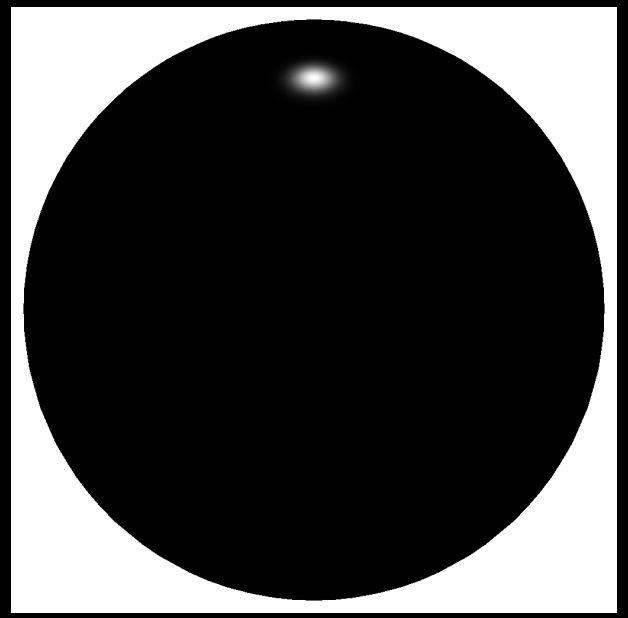

Sharpness=1000

vMF Axis

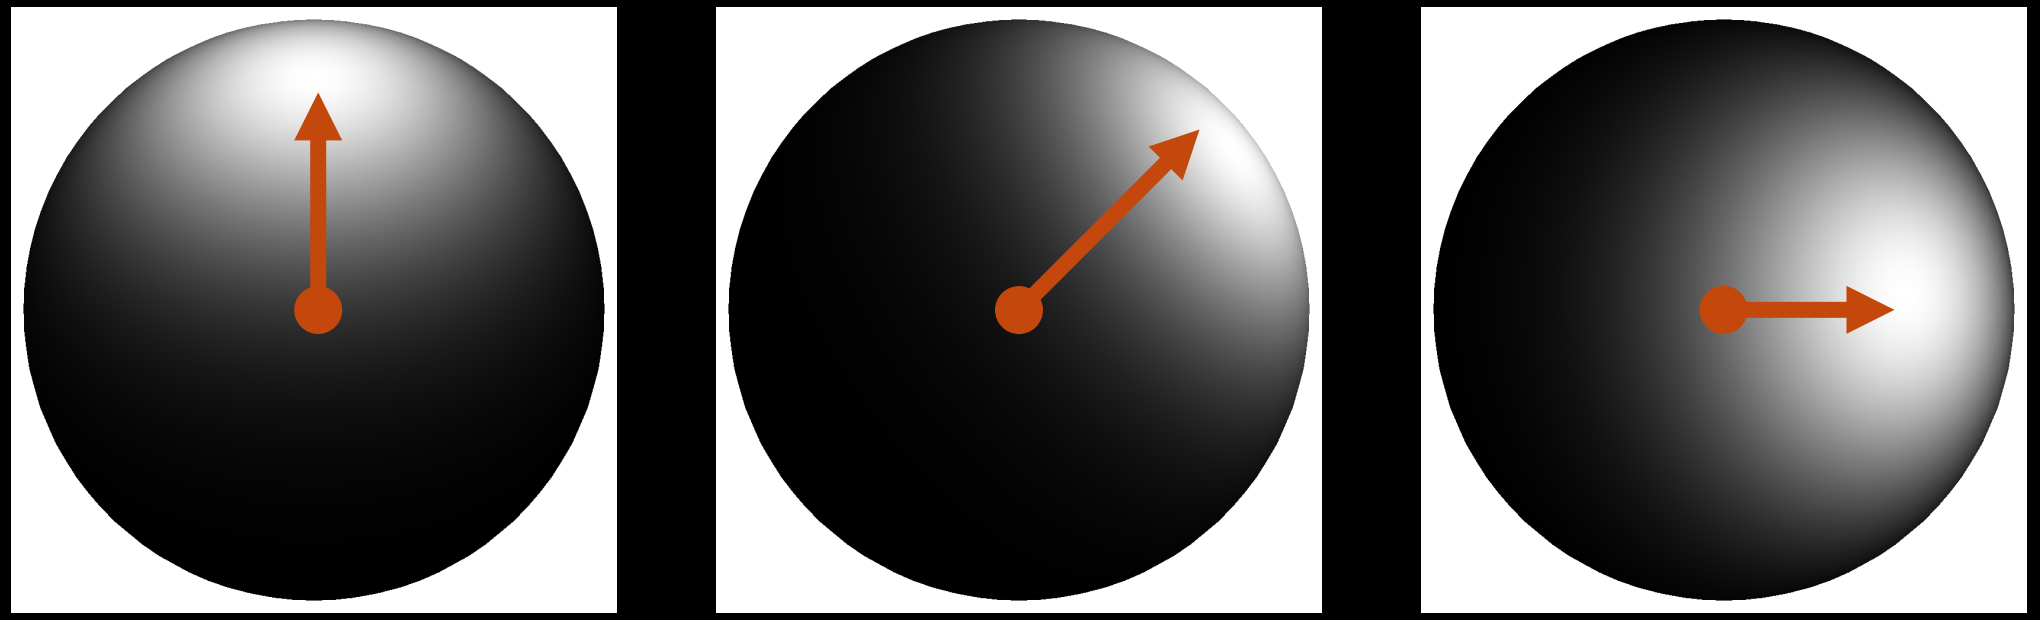

Estimation of the vMF Lobe for Each Pixel

- vMF parameters are obtained from the average direction $\acute{\mathbf{v}}_s$ of the PDF [Banerjee et al. 2005]

vMF axis = $\frac{\acute{\mathbf{v}}_s}{\|\acute{\mathbf{v}}_s\|}$

vMF sharpness = $\frac{3\|\acute{\mathbf{v}}_s\| - \|\acute{\mathbf{v}}_s\|^3}{1 - \|\acute{\mathbf{v}}_s\|}$

$$\acute{\mathbf{v}}_s = \int_{S^2} \boldsymbol{\omega} p_s(\boldsymbol{\omega}) \, \mathrm{d}\boldsymbol{\omega}$$

Estimate the average direction by sampling directions every frame

AMD

ReSTIR with Our Rejection Heuristic

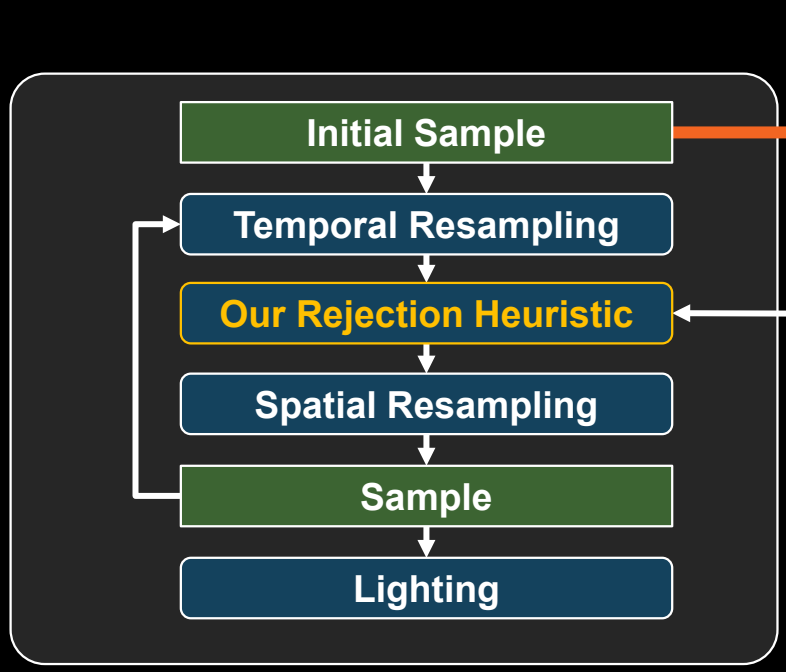

Conventional ReSTIR except for the spatial rejection heuristic

Our average-direction estimation

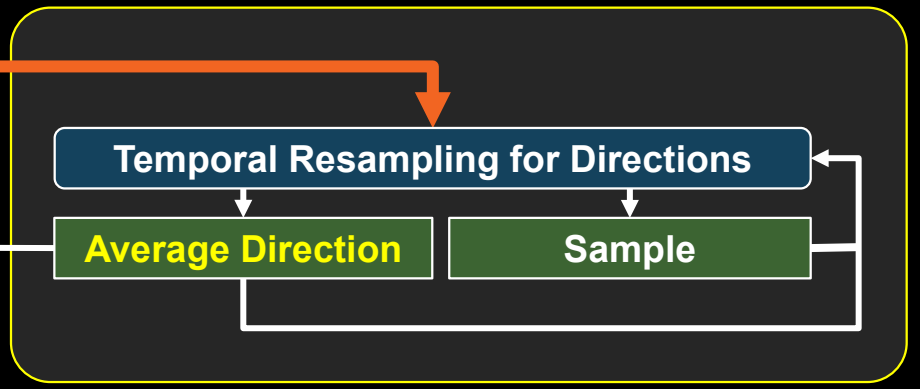

- Biased variant of ReSTIR
- Temporal resampling only
- Reuse the visibility over time [Bitterli et al. 2020]
- Reuse the initial sample from the lighting estimation
  - No additional ray tracing ☺
  - Can introduce a small bias ☹

AMD

Average Direction Estimation

- One sample is insufficient ☹
- Average the sample directions over time using resampling weights
  - A variant of *weighted importance sampling* [Bekaert et al. 2000] (or *ratio estimator* [Heitz et al. 2018])
  - Biased, but the bias reduces quickly for temporal continuities

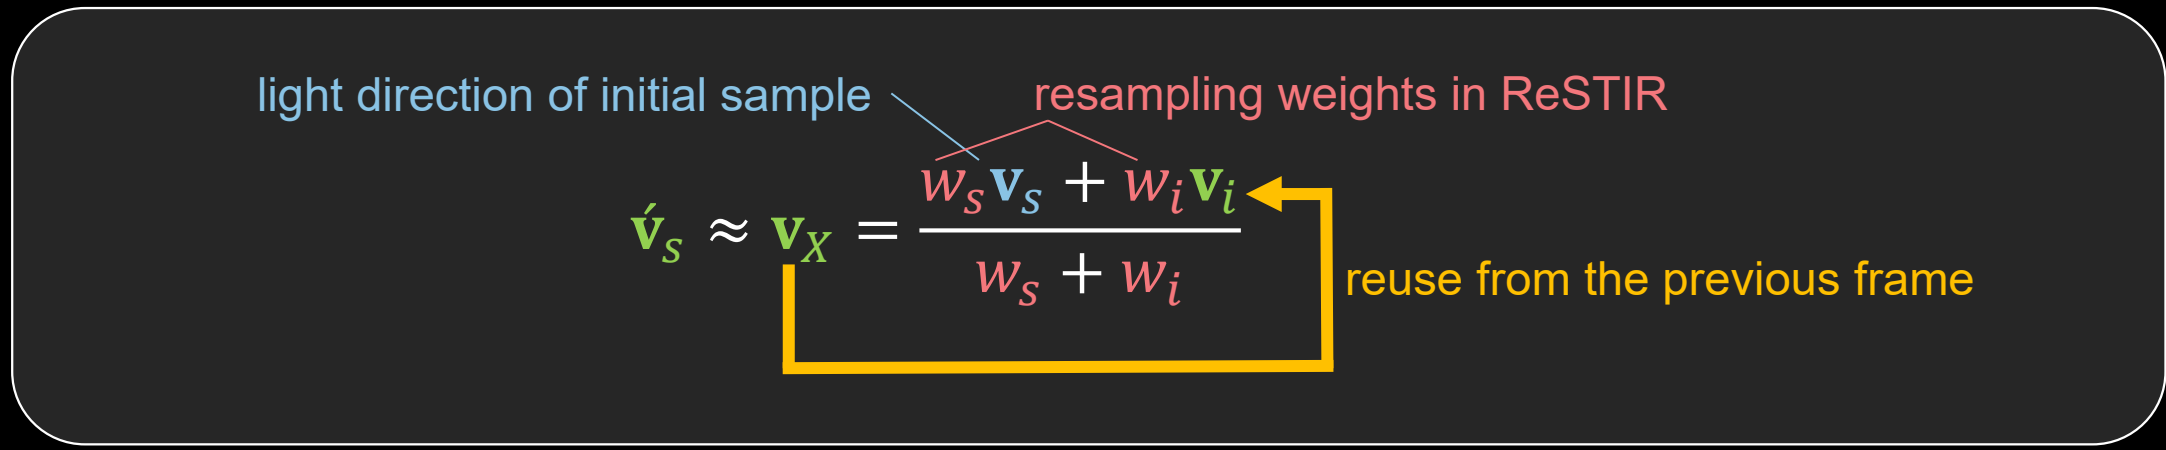

Estimated Average Directions

Rendered image

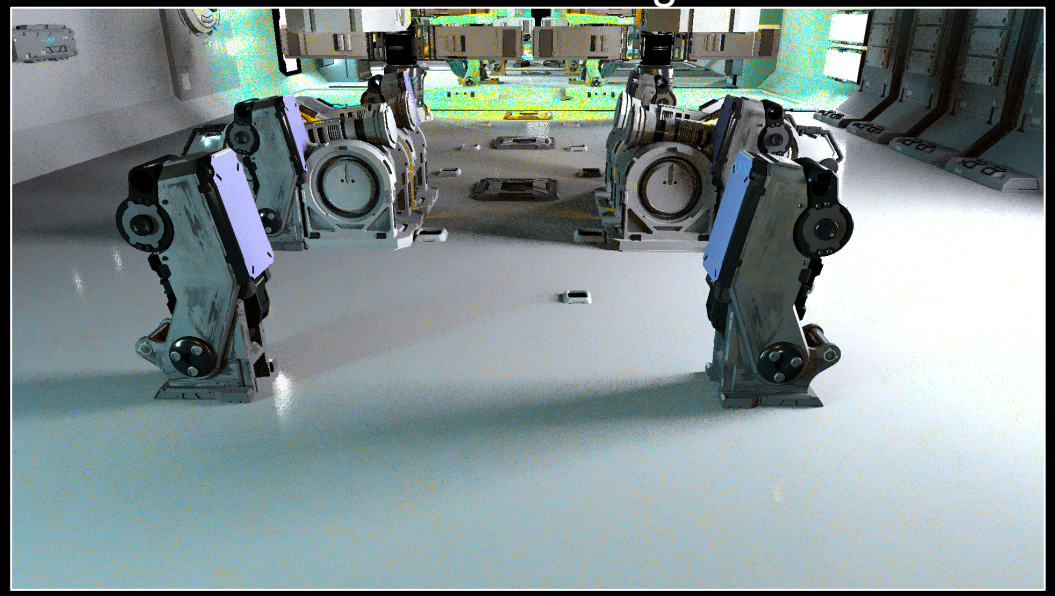

Visualization of average directions

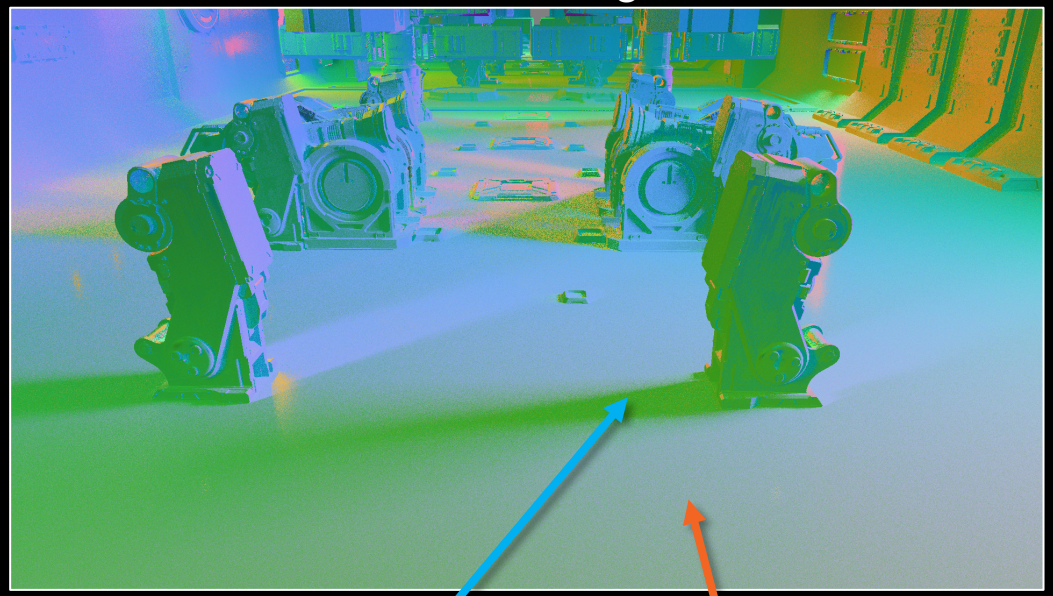

Average directions are different between shadowed and lit pixels

AMD

vMF Lobe Similarity

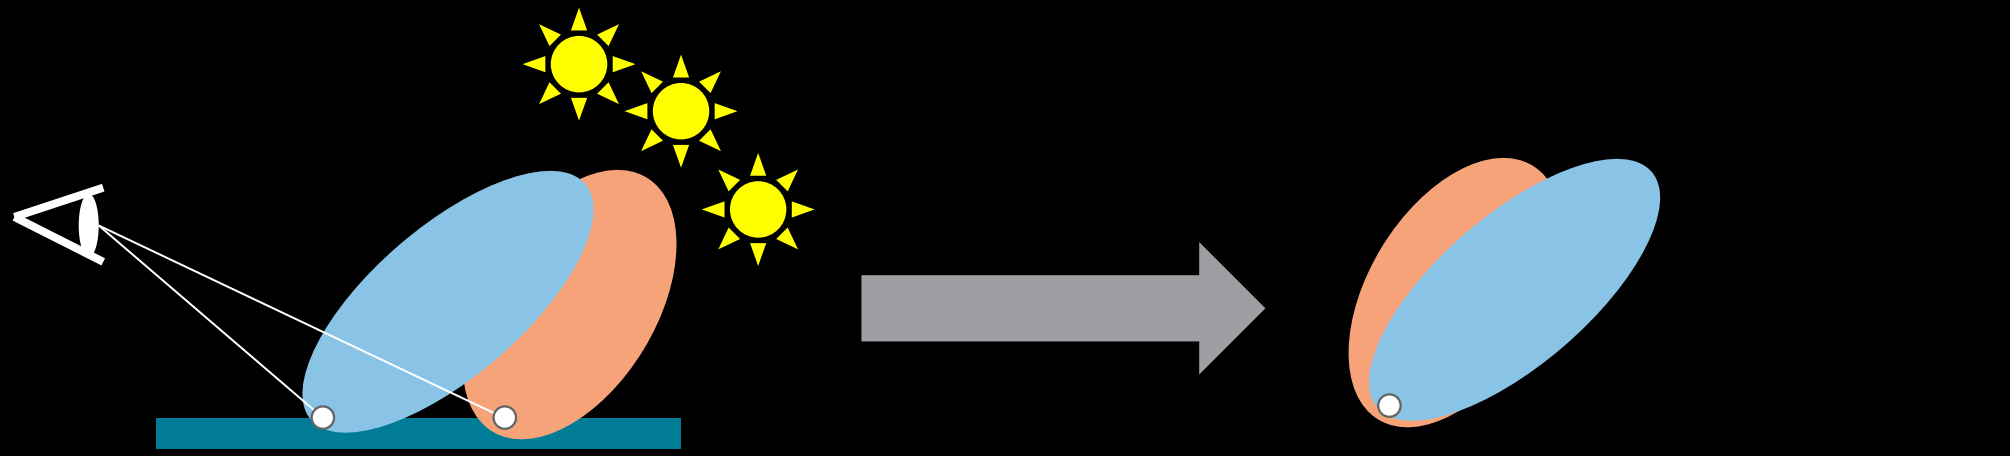

Estimated vMF lobes at two pixels

Overlap of the lobes in spherical domain

- Overlap of two vMF lobes = product integral-based similarity [Tokuyoshi 2015]
  - Simple analytical solution is available ☺
- Ignore the shift of light directions between pixels?
  - Our simple approach: Smooth the estimated vMF to consider the similarity of shifted directions
  - Please see our paper for details

AMD

Results (Visibility Reuse, 2 rpp)

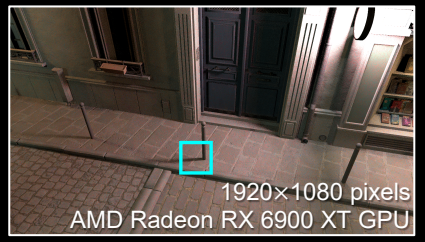

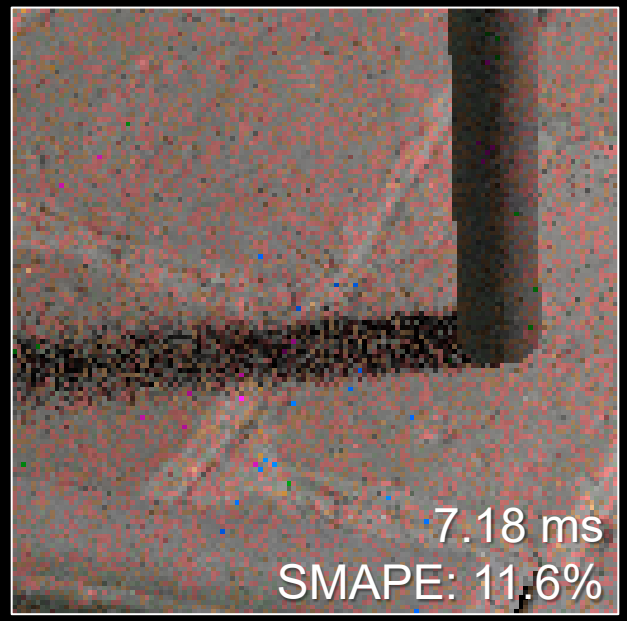

Previous method

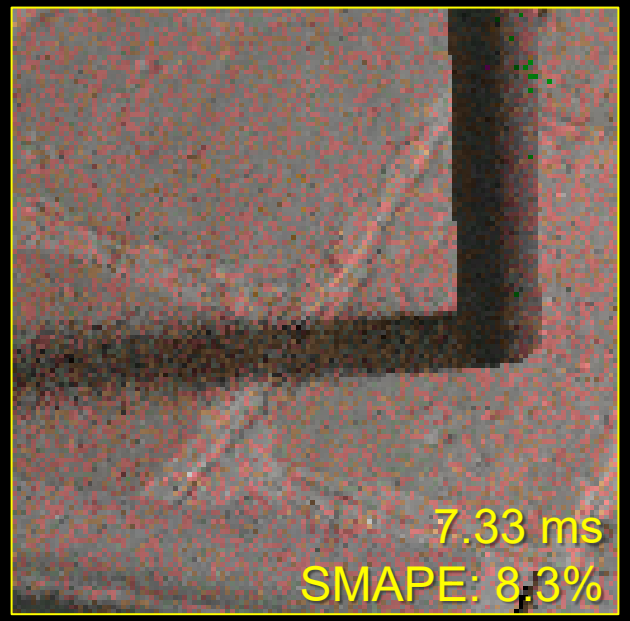

Ours

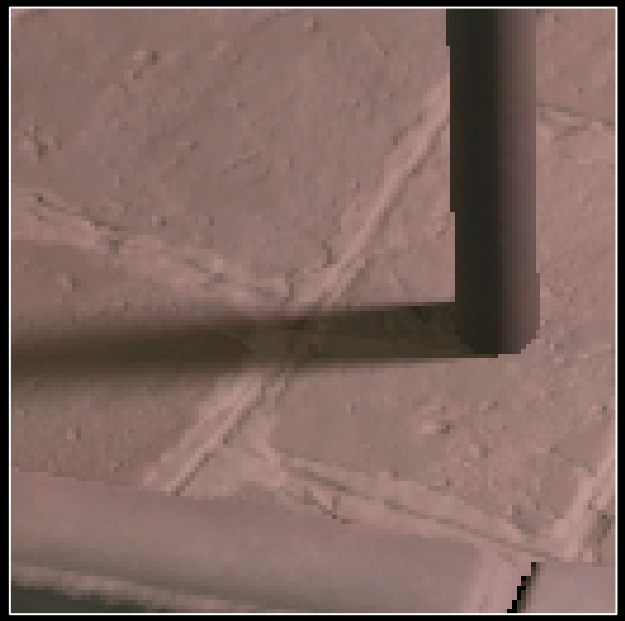

Reference

Our method reduces variance as well as darkening bias for contact shadows ☺

AMD

Results (Visibility Reuse, 1.75 rpp)

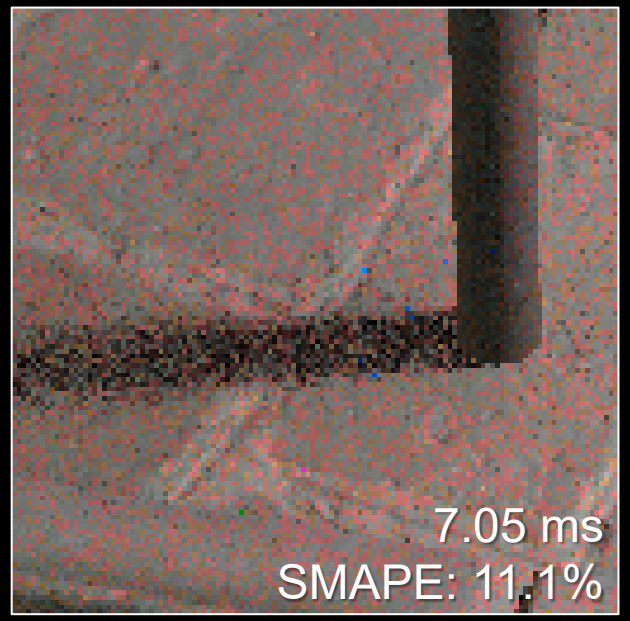

Previous method

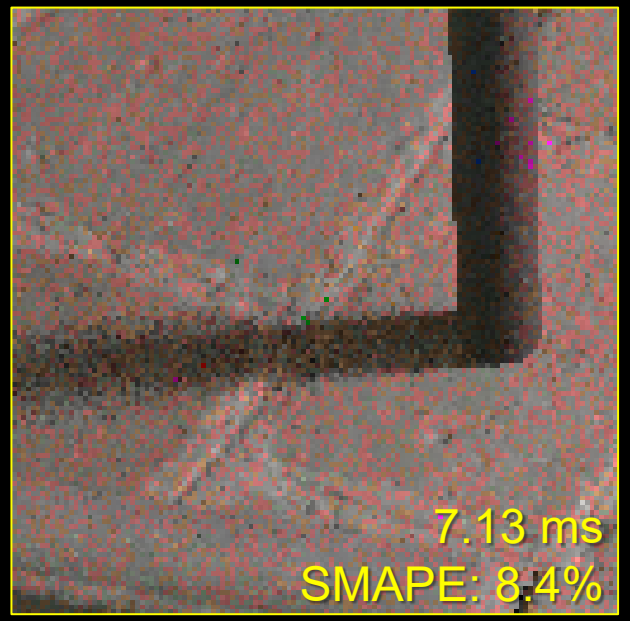

Ours

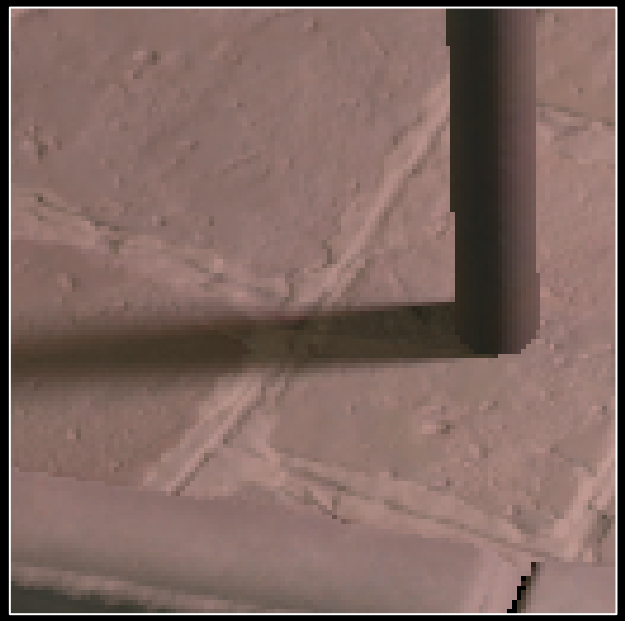

Reference

Results (Visibility Reuse, 1.5 rpp)

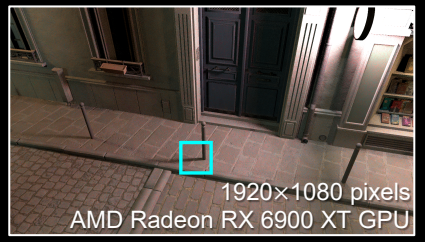

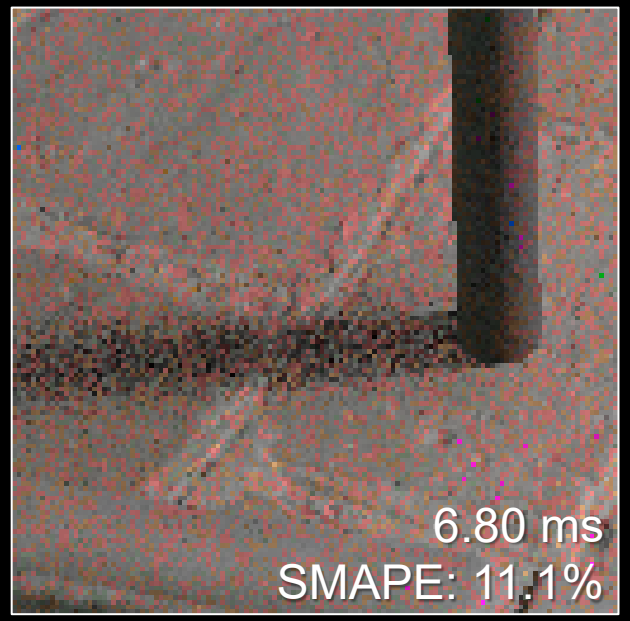

Previous method

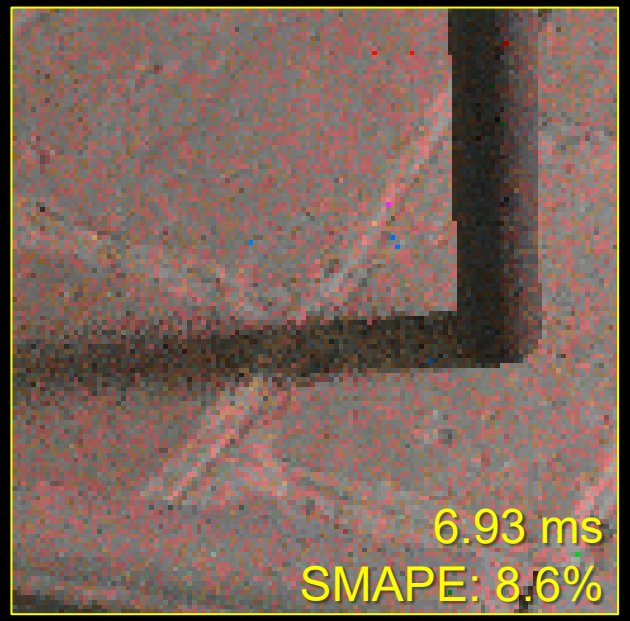

Ours

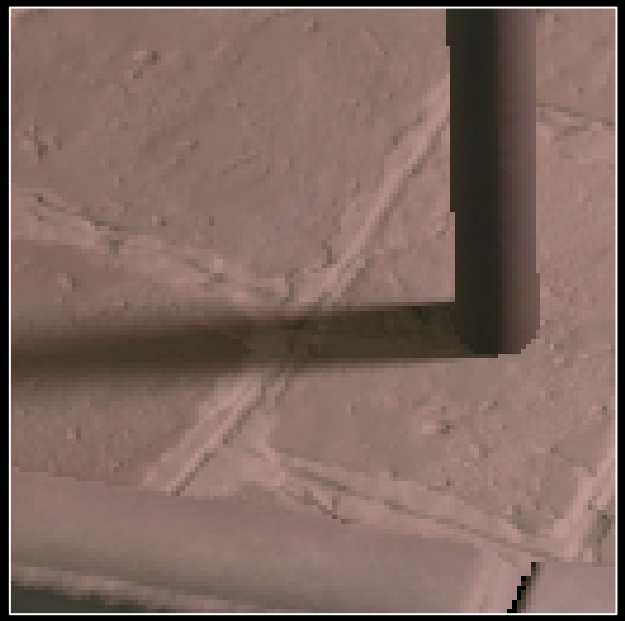

Reference

Results (Visibility Reuse, 1.25 rpp)

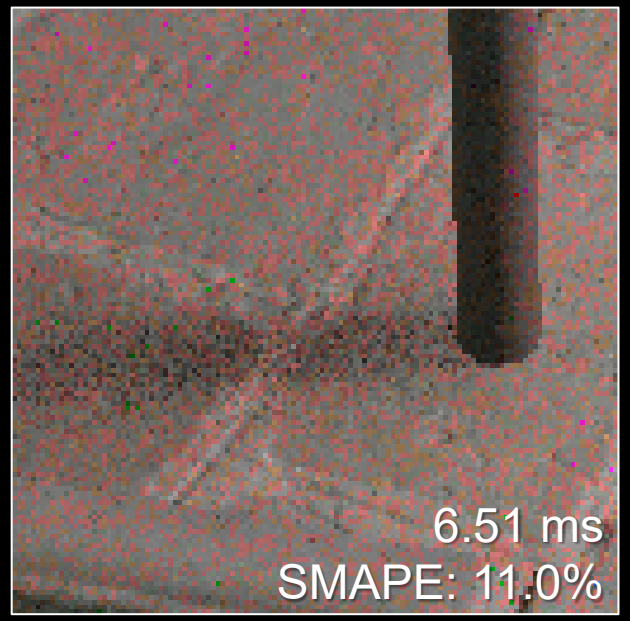

Previous method

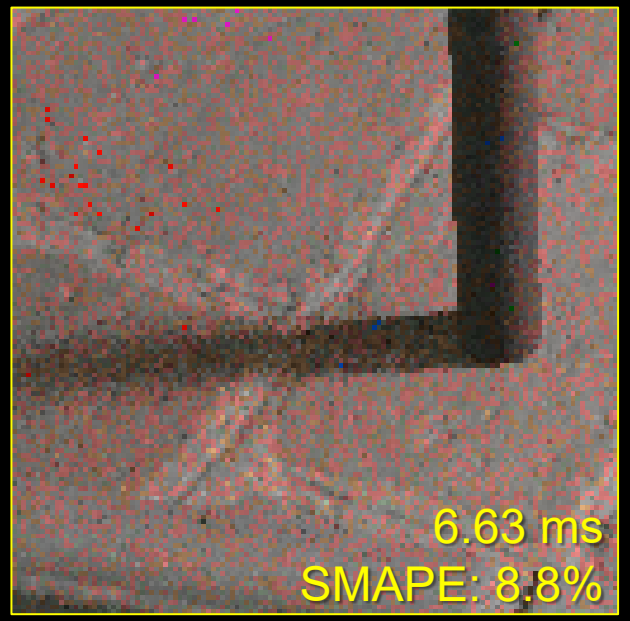

Ours

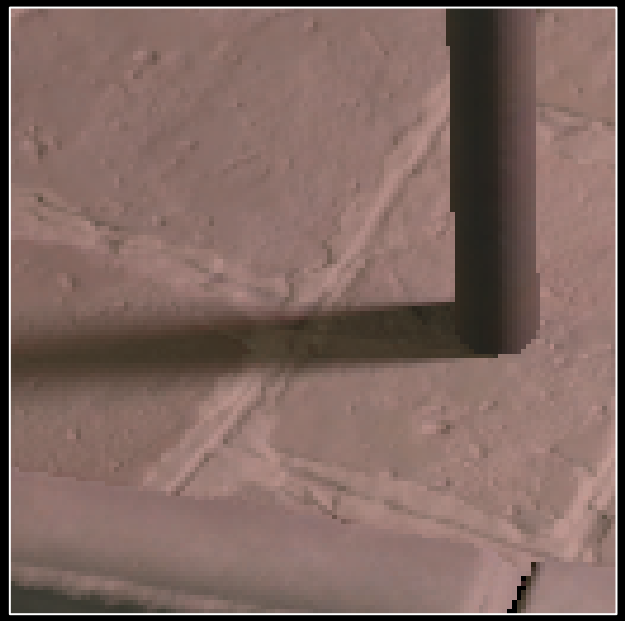

Reference

Results (Visibility Reuse, 1 rpp)

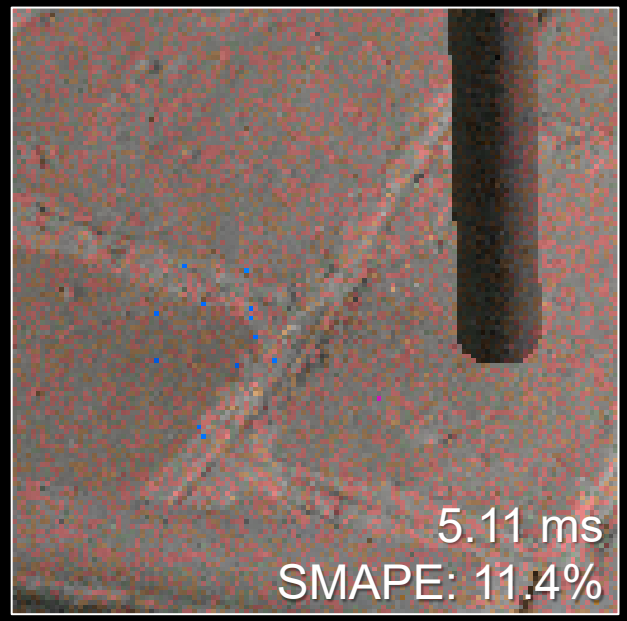

Previous method

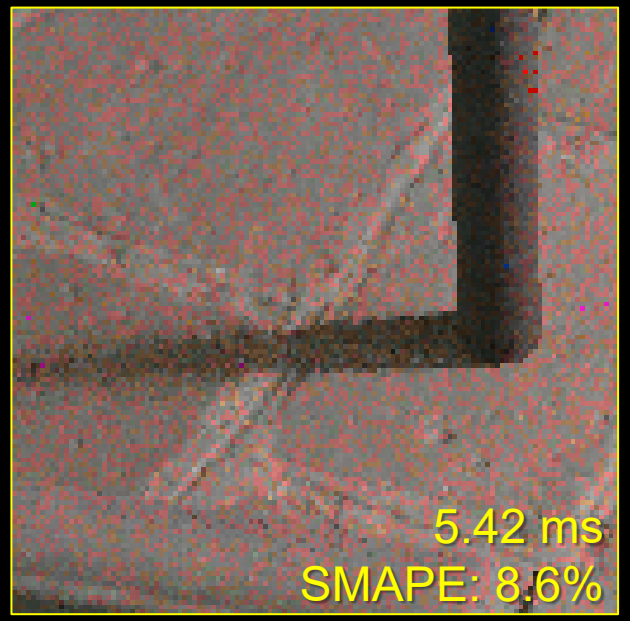

Ours

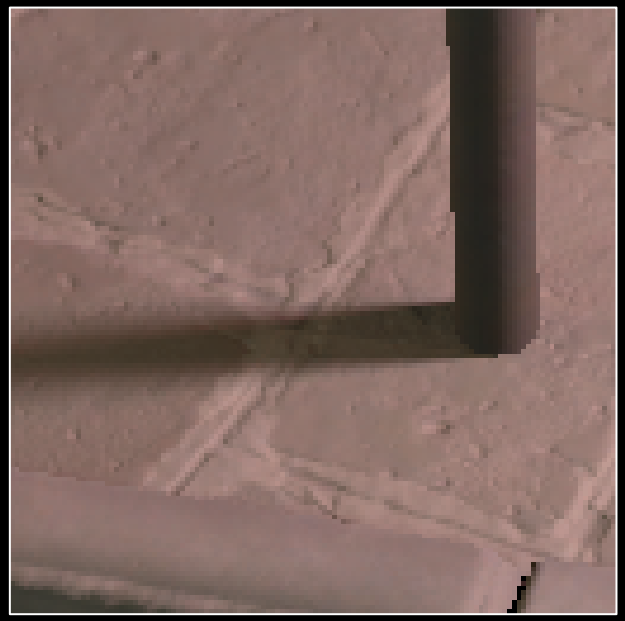

Reference

Our method preserves contact shadows more than the previous method for a small number of rays per pixel (rpp)

AMD

Variance for No Visibility Reuse (5 rpp)

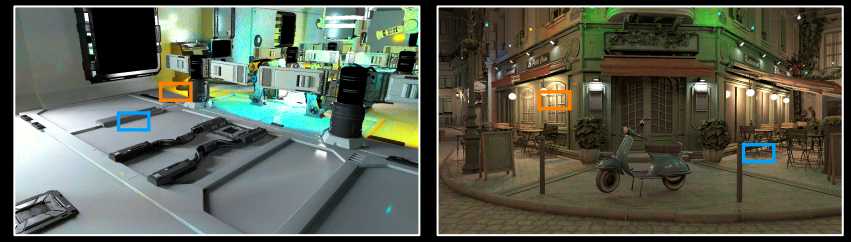

1920×1080 pixels, AMD Radeon RX 6900 XT GPU

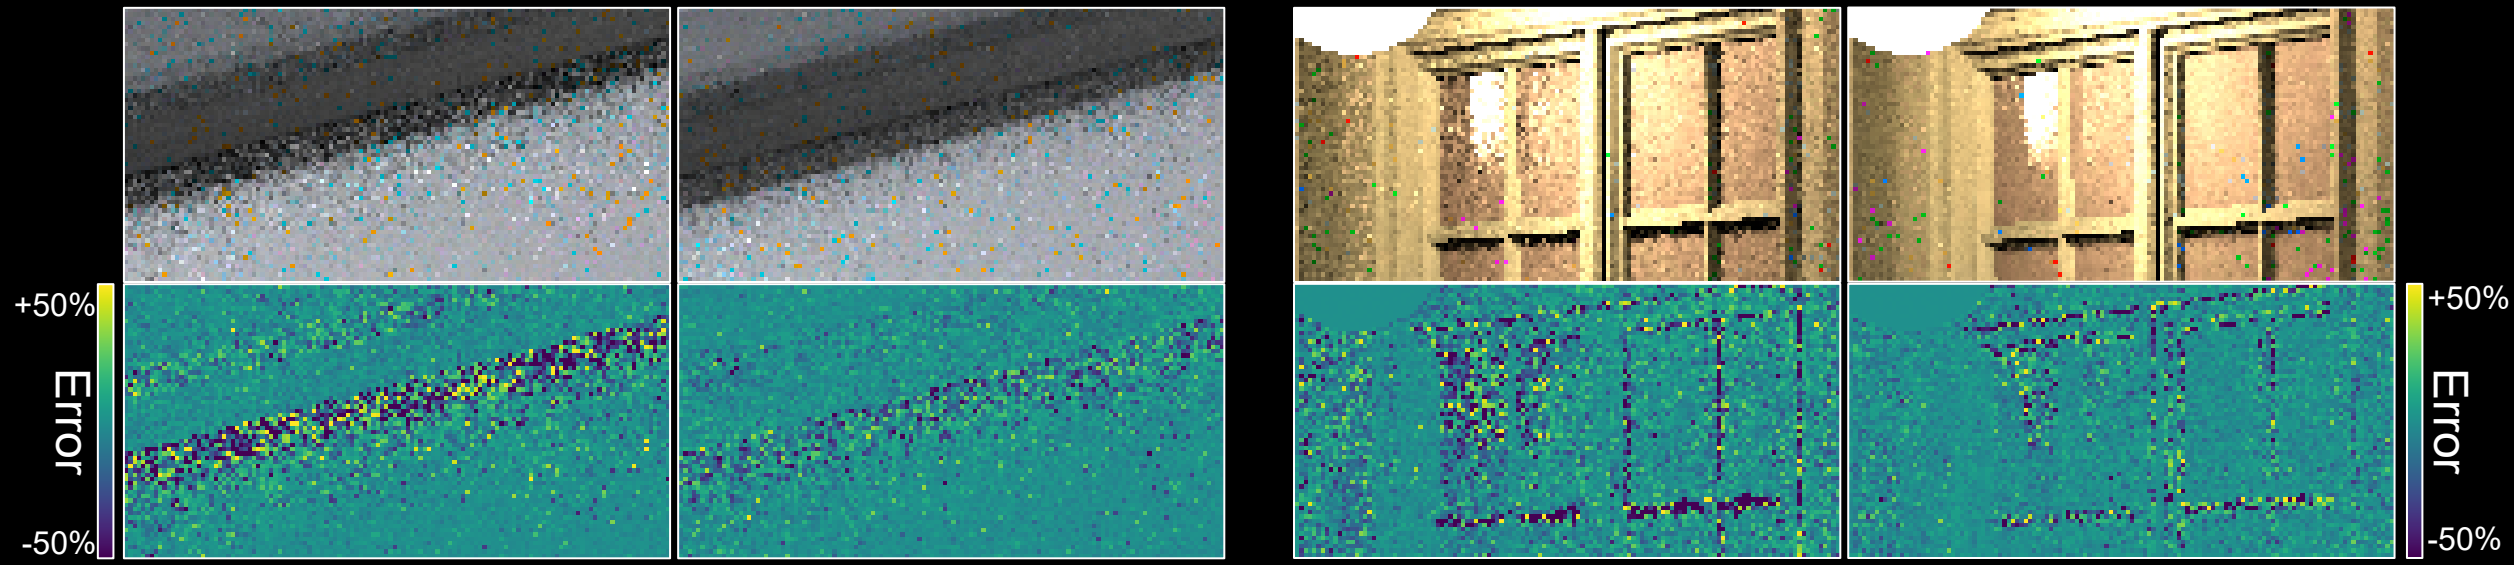

Previous
(11.7 ms, SMAPE: 7.2%)

Ours
(12.1 ms, SMAPE: 6.2%)

Previous
(19.1 ms, SMAPE: 11.1%)

Ours
(19.7 ms, SMAPE: 9.4%)

Our method reduces variance ☺

AMD

Bias for Converged Images (w/o Visibility Reuse)

Reference | Our method | Visualization of error

- Small bias in dark shadows
- Barely perceptible in the rendered image

AMD

Bias for Converged Images (w/o Visibility Reuse)

Reference

AMD

Bias for Converged Images (w/o Visibility Reuse)

Our method

AMD

Handling Temporal Estimation Error for Animation

AMD

Estimation Error of the Average Direction

- Our heuristic can fail for animation
- Significant bias & variance when the accumulated sample count is small
  - E.g., temporal disocclusions
- Variance of our PDF similarity can correlate to the variance of lighting
  - Average-direction estimation shares the initial sample with lighting estimation
  - This correlation results in a bias in the rejection of spatial reuse
  - Decorrelate using different random numbers for every resampling routine
    - Still noticeable for temporal disocclusions
- Temporal delay due to averaging samples over time

AMD

Combination with Existing Heuristics

- Use our rejection heuristic only when the sample count is sufficient
- Combine our heuristic and existing heuristics using the temporally accumulated sample count

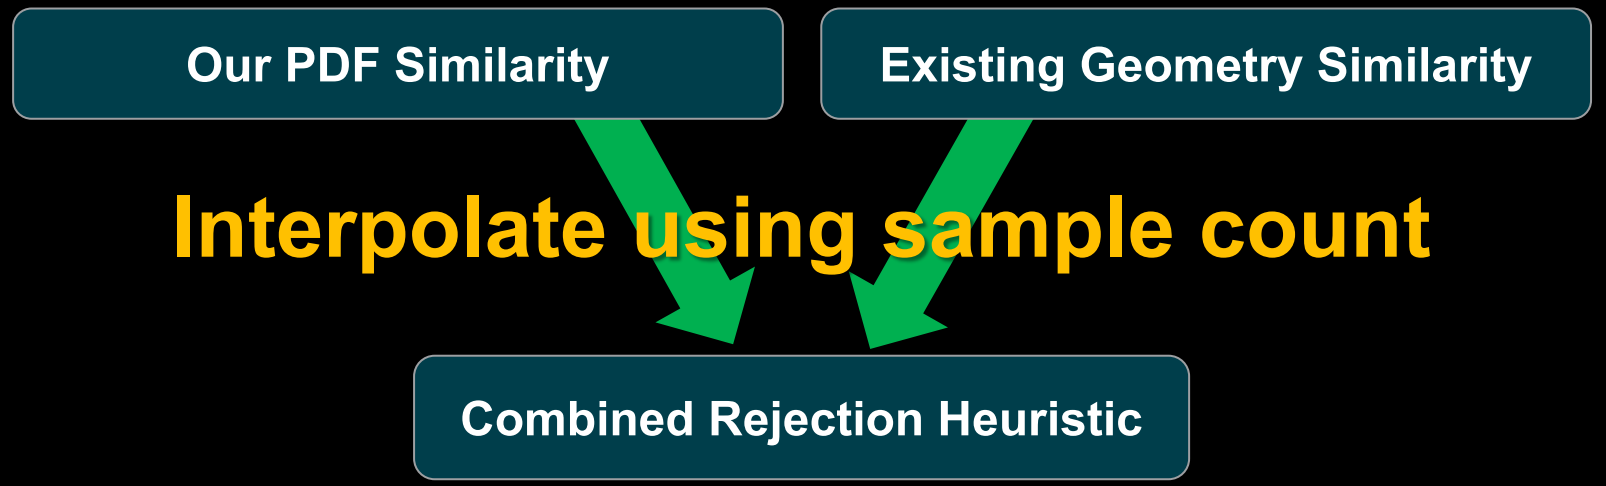

Our heuristic is effective only for temporally continuous pixels

AMD

Reduction of Delay for Moving Lights

- Reduce the accumulated sample count for the average-direction estimation when a reused sample light moves
  - Introduce a bias in the average direction estimation, but not a problem
  - Existing heuristics become more dominant than our heuristic for this case

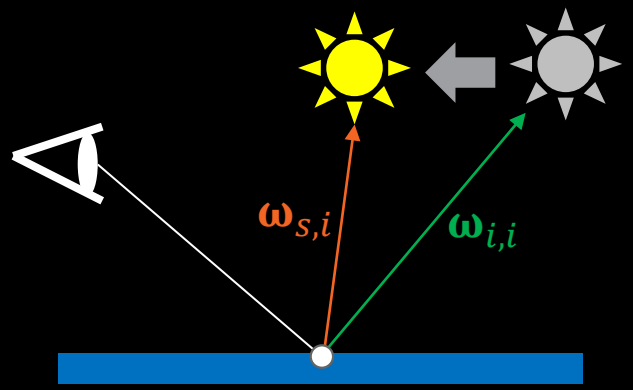

Spherical Gaussian-based reduction rate

$$\exp\left(\lambda\left(\omega_{s,i} \cdot \omega_{i,i}\right) - 1\right)$$

User-specified constant to control the sensitivity of moving lights

AMD

Results

1920×1080 PIXELS

1 M VPLS GENERATED ON AREA LIGHTS

AMD RADEON RX 6900 XT GPU

AMD

Visibility-Reuse ReSTIR with 1 rpp

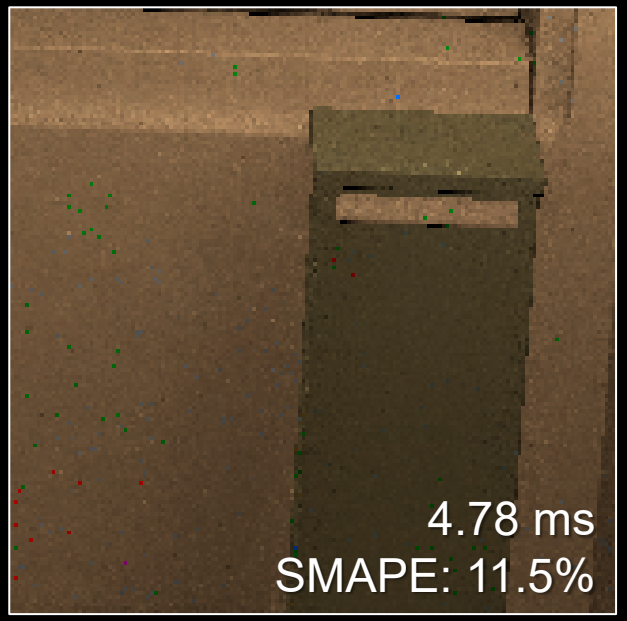

Previous

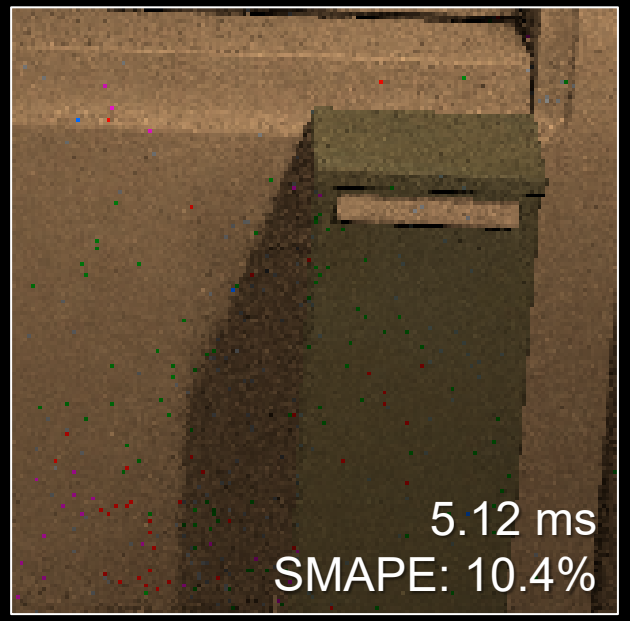

Ours

Ours in motion

- Can render shadow edges for temporal continuous pixels
- Lose shadows in motion
  - Our method uses the previous heuristic for temporal discontinuities

AMD

Visibility-Reuse ReSTIR with Adaptive Ray Tracing

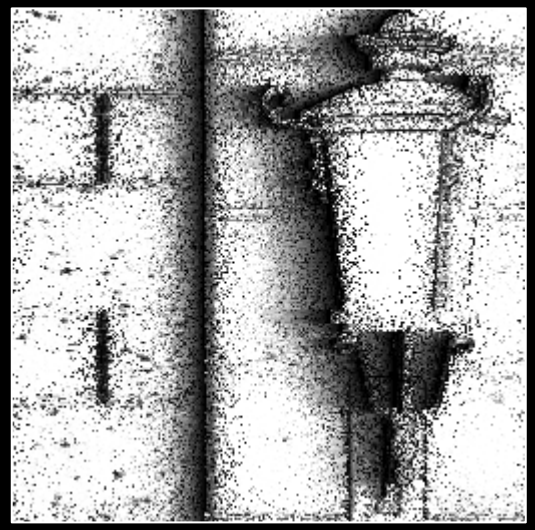

temporal continuity
based on # of dir. samples

Spatially adaptive
(1.25 rpp, SMAPE: 25.2%)

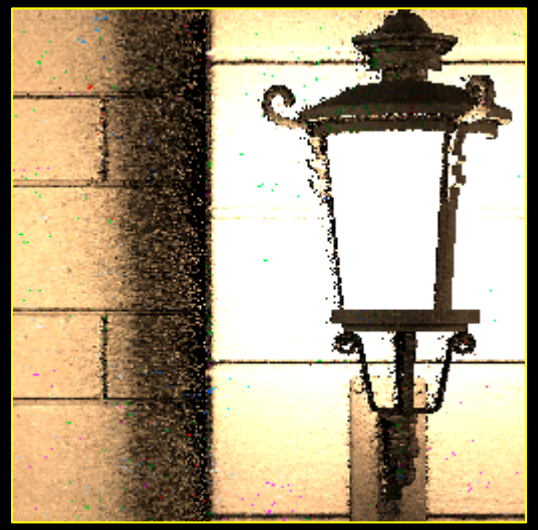

Spatiotemporally adaptive
(1.53 rpp, SMAPE: 26.6%)

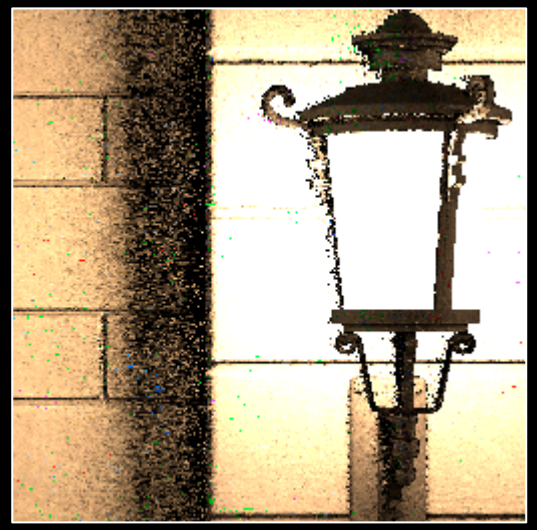

Non-adaptive
(2 rpp, SMAPE: 27.4%)

Spatiotemporally adaptive ray tracing compensates shadows for temporal disocclusions

AMD


CPU 0.123 ms, GPU 11.705 ms, 60 Hz
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Ours (adaptive) ReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
Restaurant
Le Petit Cafe
Our method with adaptive ray tracing for temporal discontinuities
It renders highly detailed and temporally coherent shadows using less than two rays

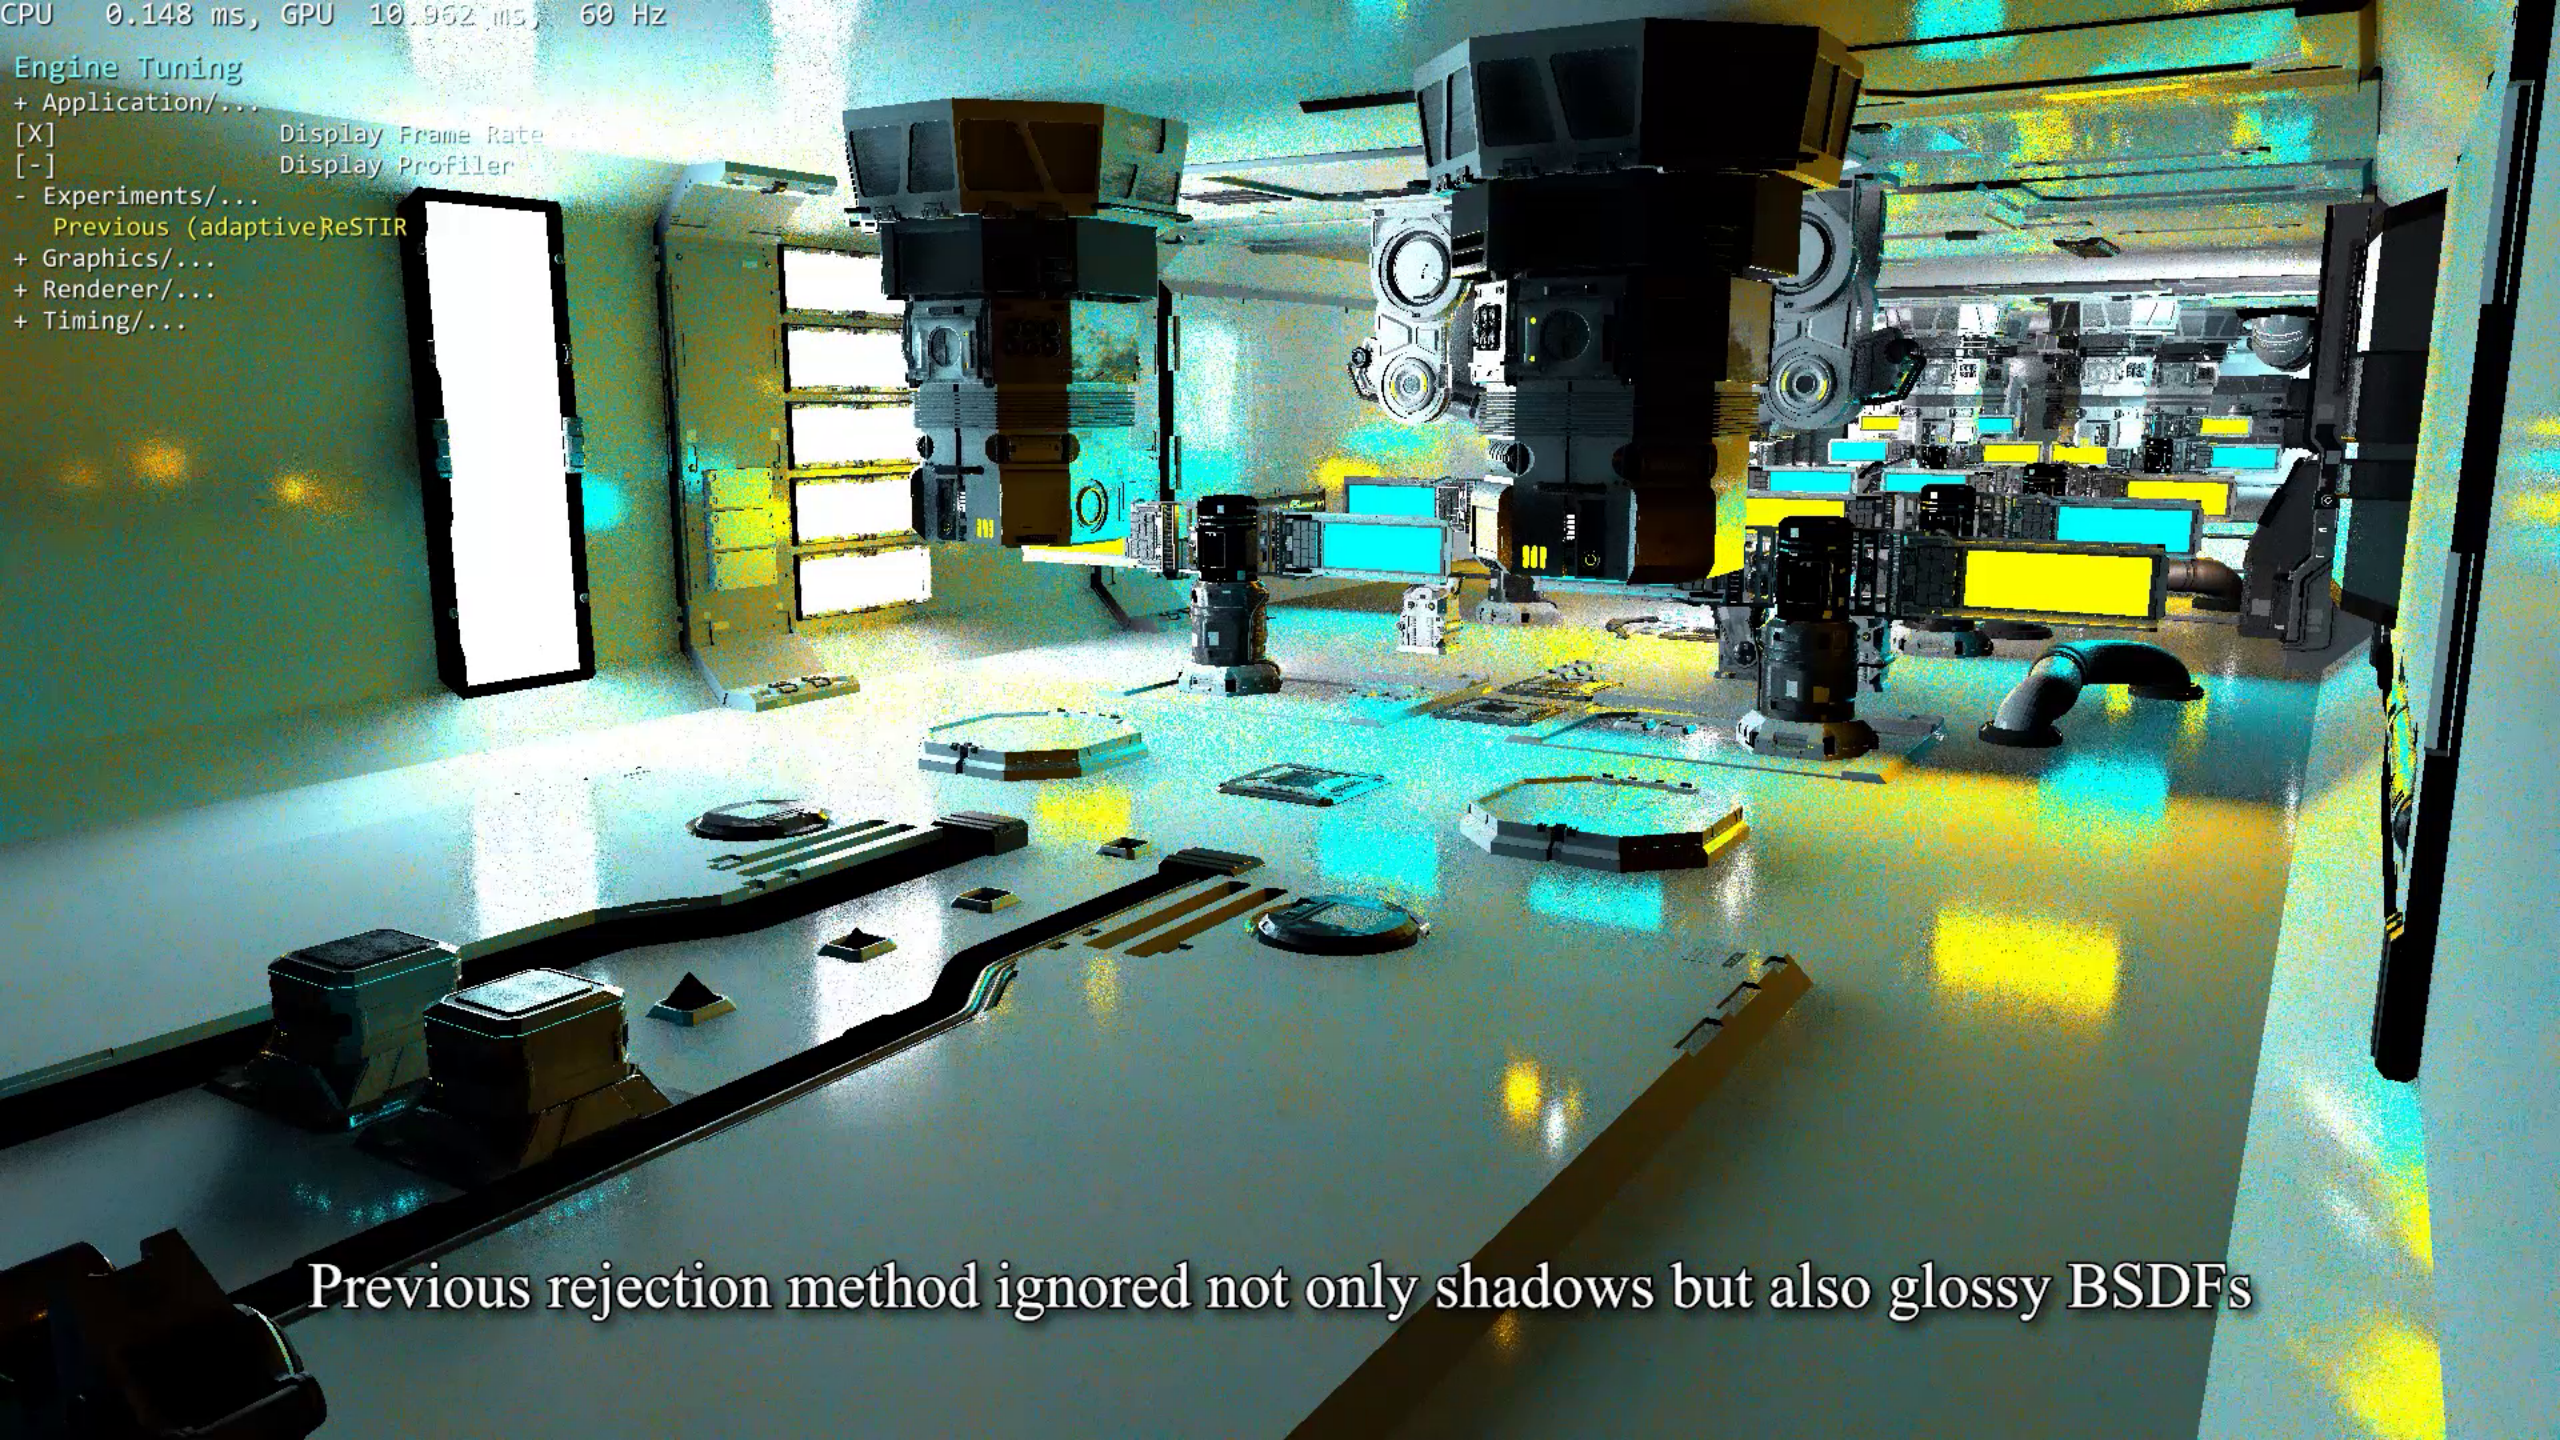
CPU 0.148 ms, GPU 10.962 ms, 60 Hz
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Previous (adaptiveReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
Previous rejection method ignored not only shadows but also glossy BSDFs

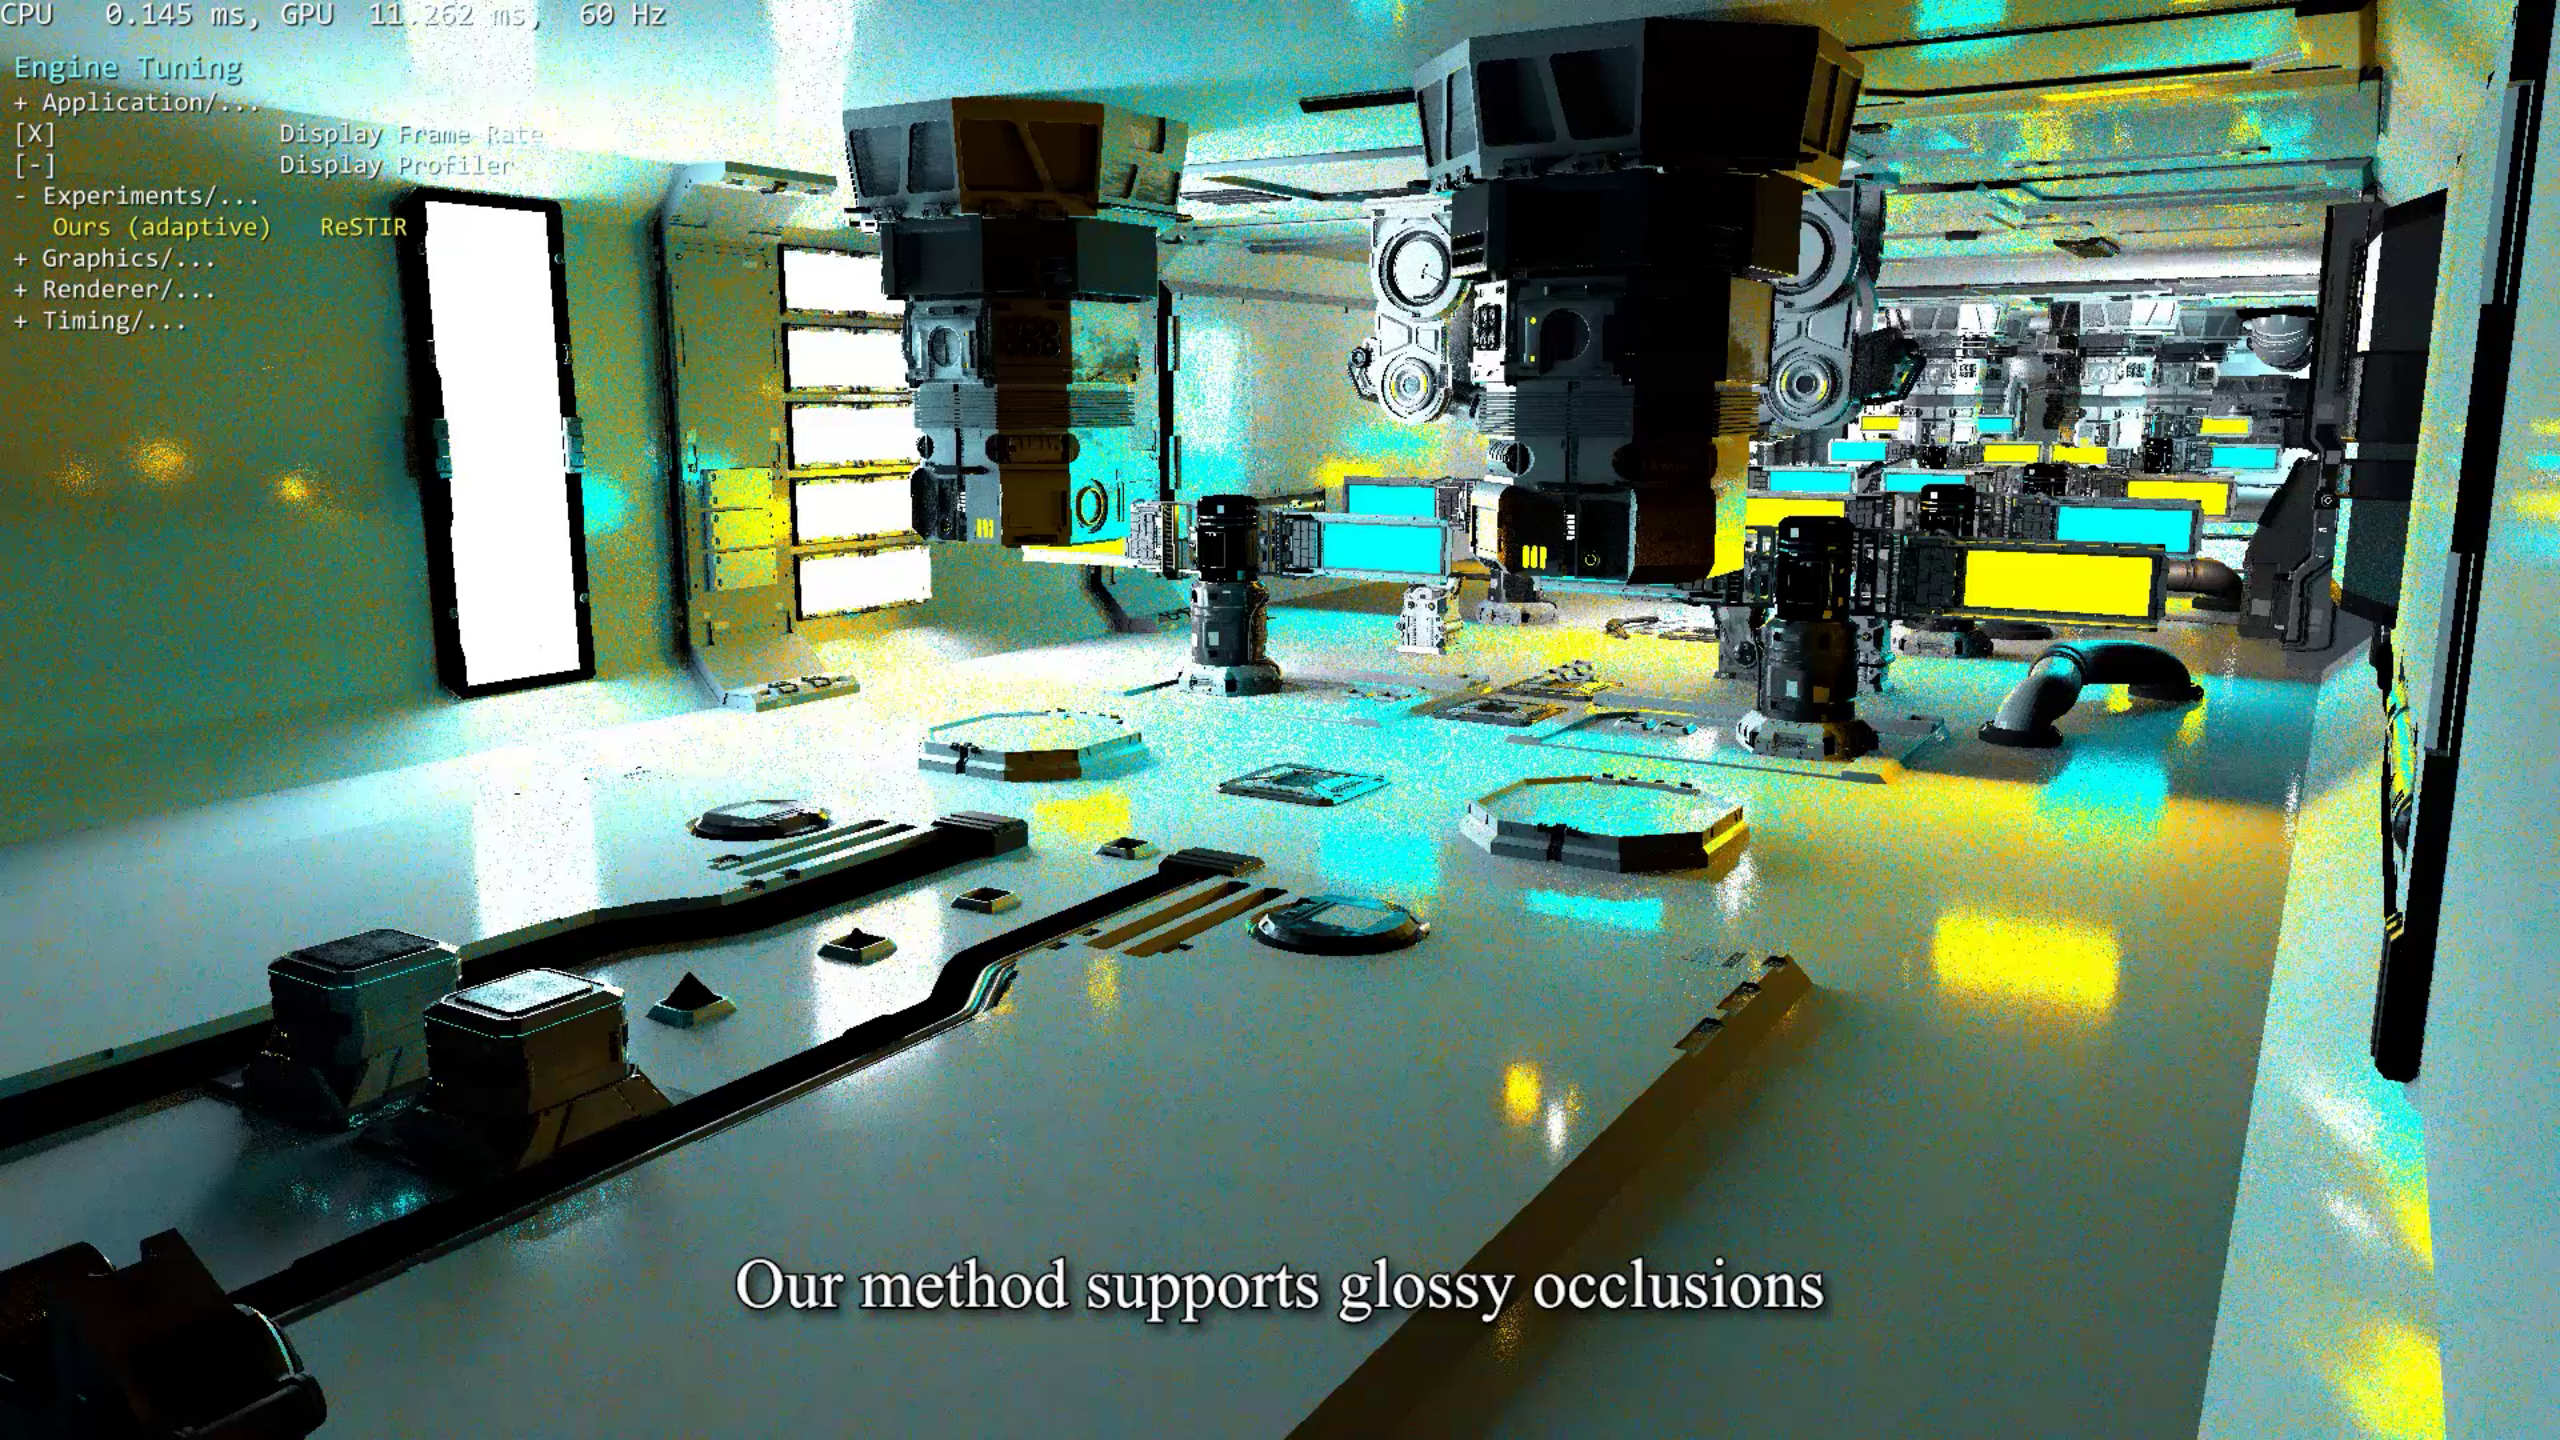
CPU 0.145 ms, GPU 11.262 ms, 60 Hz
Engine Tuning
+ Application/...
[X] Display Frame Rate
[-] Display Profiler
- Experiments/...
Ours (adaptive) ReSTIR
+ Graphics/...
+ Renderer/...
+ Timing/...
Our method supports glossy occlusions

Limitations

- Biased & inconsistent estimator
  - The bias is negligibly small for temporally continuous pixels
- Our rejection heuristic works only for temporal continuities
  - Use the previous method as a fall back for temporal discontinuities
  - Temporally adaptive shadow ray tracing can compensate lost shadows
- Our method can approximate different PDFs into an identical vMF lobes
  - False positives
- Single-bounce only
- Memory overhead for average-direction estimation
  - 16 bytes/pixel for a reservoir in our experimental implementation

AMD

Limitations (cont'd)

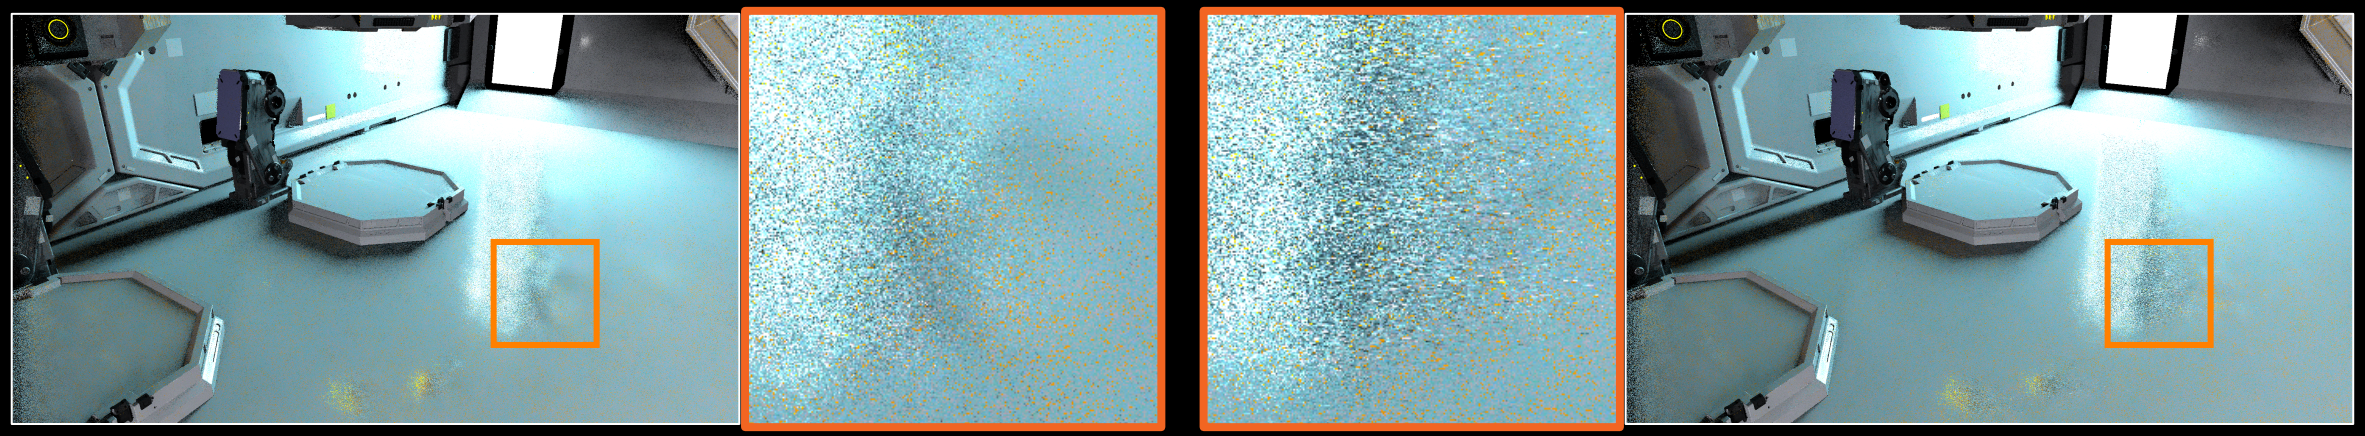

Previous method (7.52 ms, SMAPE: 17.3%)

Ours (7.72 ms, SMAPE: 16.9%)

- Our PDF similarity estimation has a temporal delay for highly glossy surfaces viewed from moving camera
  - Variance for moving highlights

- Future work:
  - Glossy lobe similarity [Tokuyoshi 2015]
  - Decoupling of incoming radiance and BSDFs [Wang et al. 2009]

AMD

Conclusion

- Rejection heuristic based on the similarity of PDF shapes
- Efficient vMF approximation & temporal estimation for the PDF shape
- Stable combination of our heuristic and previous heuristics
  - Comparable quality to the existing heuristic for temporal disocclusions
- Improve the image quality for temporal continuities

AMD

References

- Arindam Banerjee, Inderjit S. Dhillon, Joydeep Ghosh, and Suvrit Sra. 2005. Clustering on the Unit Hypersphere Using von Mises-Fisher Distributions. J. Mach. Learn. Res. 6 (2005), 1345–1382
- Philippe Bekaert, Mateu Sbert, and Yves D. Willems. 2000. Weighted Importance Sampling Techniques for Monte Carlo Radiosity. In EGWR ’00. 35–46
- Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal Reservoir Resampling for Real-Time Ray Tracing with Dynamic Direct Lighting. ACM Trans. Graph. 39, 4, Article 148 (2020), 17 pages
- Ronald Aylmer Fisher. 1953. Dispersion on a sphere. Proc. R. Soc. Lond. Ser. A 217, 1130 (1953), 295–305
- Eric Heitz, Stephen Hill, and Morgan McGuire. 2018. Combining Analytic Direct Illumination and Stochastic Shadows. In I3D ’18. Article 2, 11 pages
- Yusuke Tokuyoshi. 2015. Specular Lobe-Aware Filtering and Upsampling for Interactive Indirect Illumination. Comput. Graph. Forum 34, 6 (2015), 135–147
- Yu-Ting Tsai and Zen-Chung Shih. 2006. All-Frequency Precomputed Radiance Transfer Using Spherical Radial Basis Functions and Clustered Tensor Approximation. ACM Trans. Graph. 25, 3 (2006), 967–976
- Jiaping Wang, Peiran Ren, Minmin Gong, John Snyder, and Baining Guo. 2009. All-Frequency Rendering of Dynamic, Spatially-Varying Reflectance. ACM Trans. Graph. 28, 5 (2009), 133:1–133:10
- Chris Wyman and Alexey Panteleev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In HPG ’21. 19 pages

AMD