Team 1

Parameter-space ReSTIR for Differentiable and Inverse Rendering
(ReSTIR DR)

Niklas Sanden
Tan Chao

Review of Team 2

SinNeRF: Training Neural Radiance Fields on Complex Scenes from a Single Image

introduce and propagate geometry pseudo labels and semantic pseudo labels to guide the progressive training process

- train this semi-supervised framework via ground truth color and depth labels of the reference view and pseudo labels on unseen views
- use image warping to obtain geometry pseudo labels and utilize adversarial training as well as a pre-trained ViT for semantic pseudo labels.

Zero-1-to-3: Zero-shot One Image to 3D Object

learn to control the camera perspective in large-scale diffusion models, enabling zero-shot novel view synthesis and 3D reconstruction from a single image.

- capitalize on the geometric priors that large-scale diffusion models learn about natural images.
- Uses a synthetic dataset to learn camera viewpoint controls, enabling generation of new images of the same object from specified angles.

SIGGRAPH 2023
LOS ANGELES+ 6-10 AUG

THE PREMIER CONFERENCE & EXHIBITION ON
COMPUTER GRAPHICS & INTERACTIVE TECHNIQUES

Parameter-space ReSTIR for Differentiable and Inverse Rendering

Wesley Chang[1], Venkataram Sivaram[1],
Derek Nowrouzezahrai[2], Toshiya Hachisuka[3],
Ravi Ramamoorthi[1], and Tzu-Mao Li[1].

[1]UCSD [2]McGill University [3]University of Waterloo

Forward and Inverse Rendering

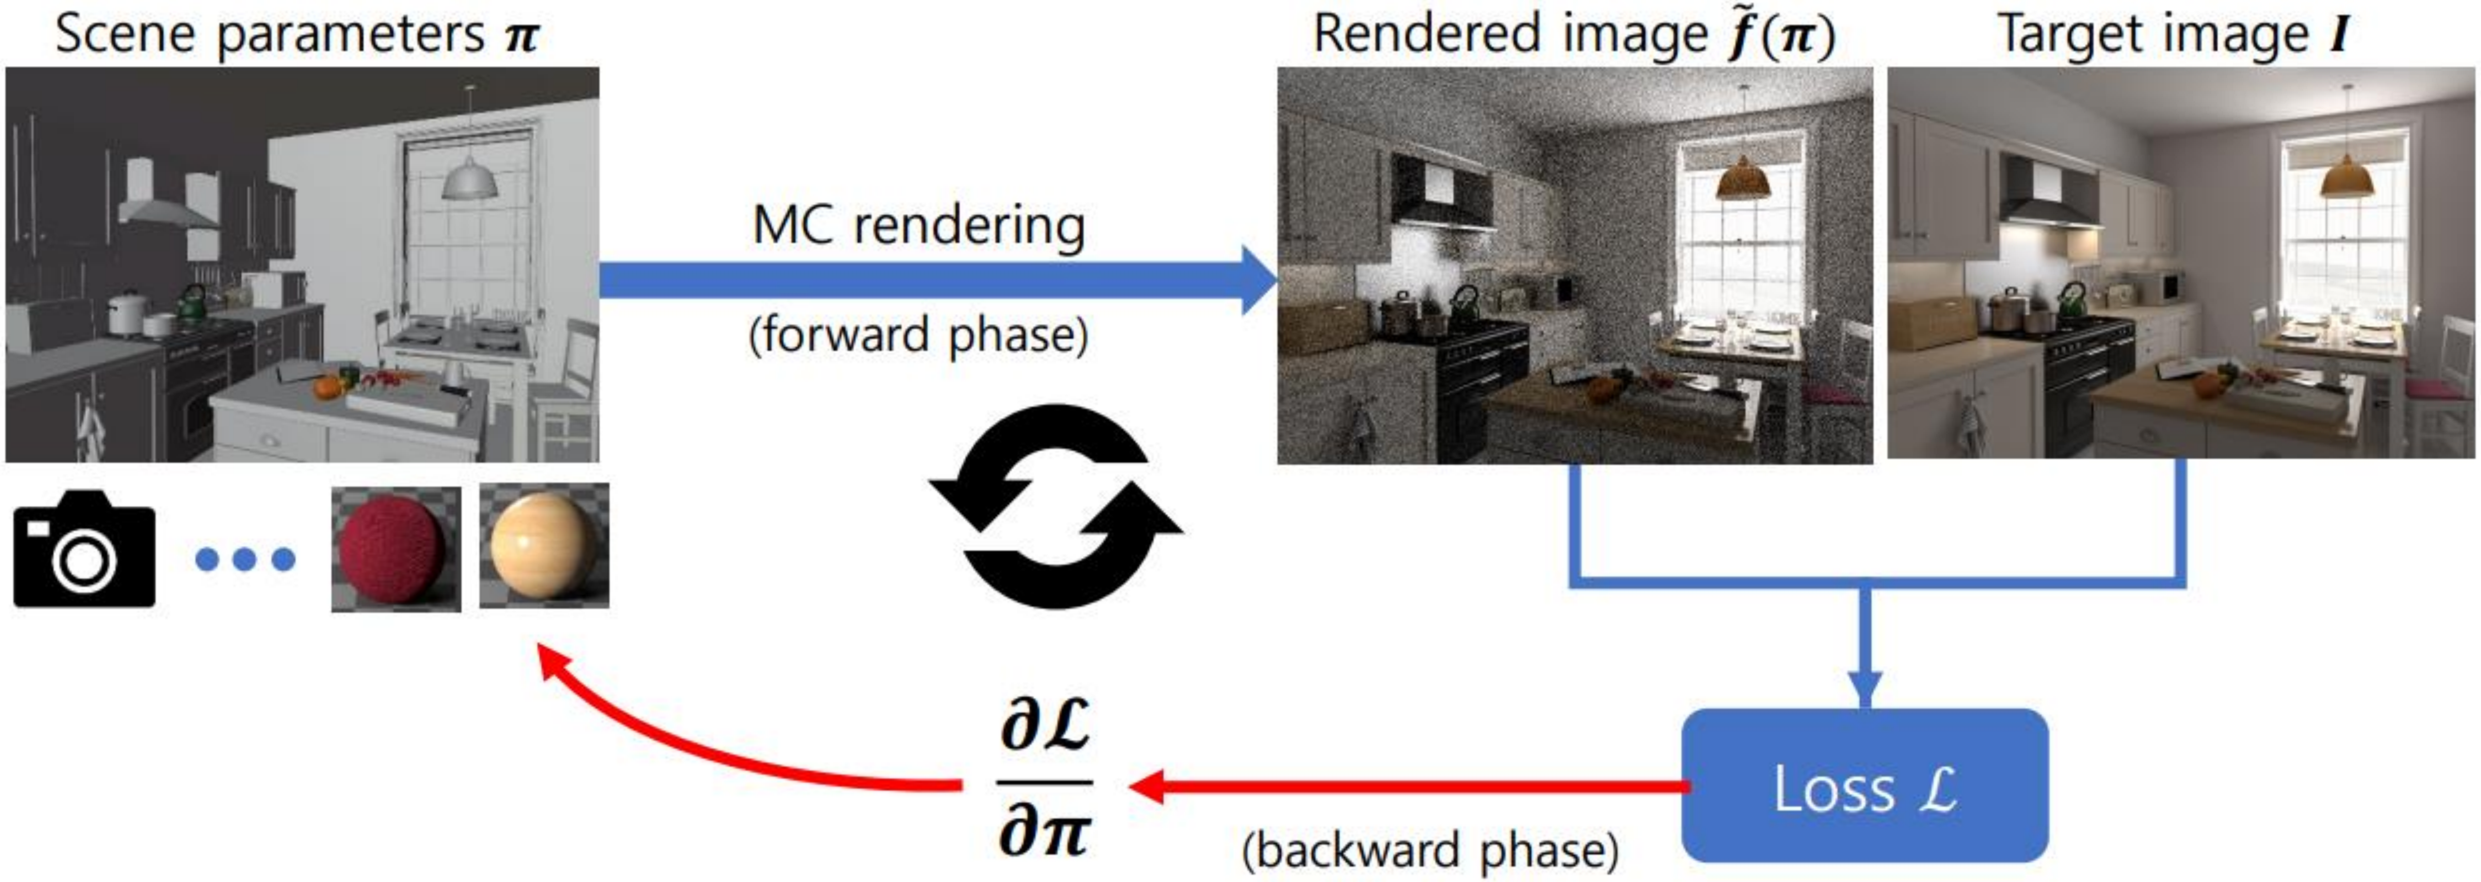

Target-Aware Image Denoising for Inverse Monte Carlo Rendering, Bochang Moon

Inverse rendering (materials)

Initial

Target

Noisy gradients

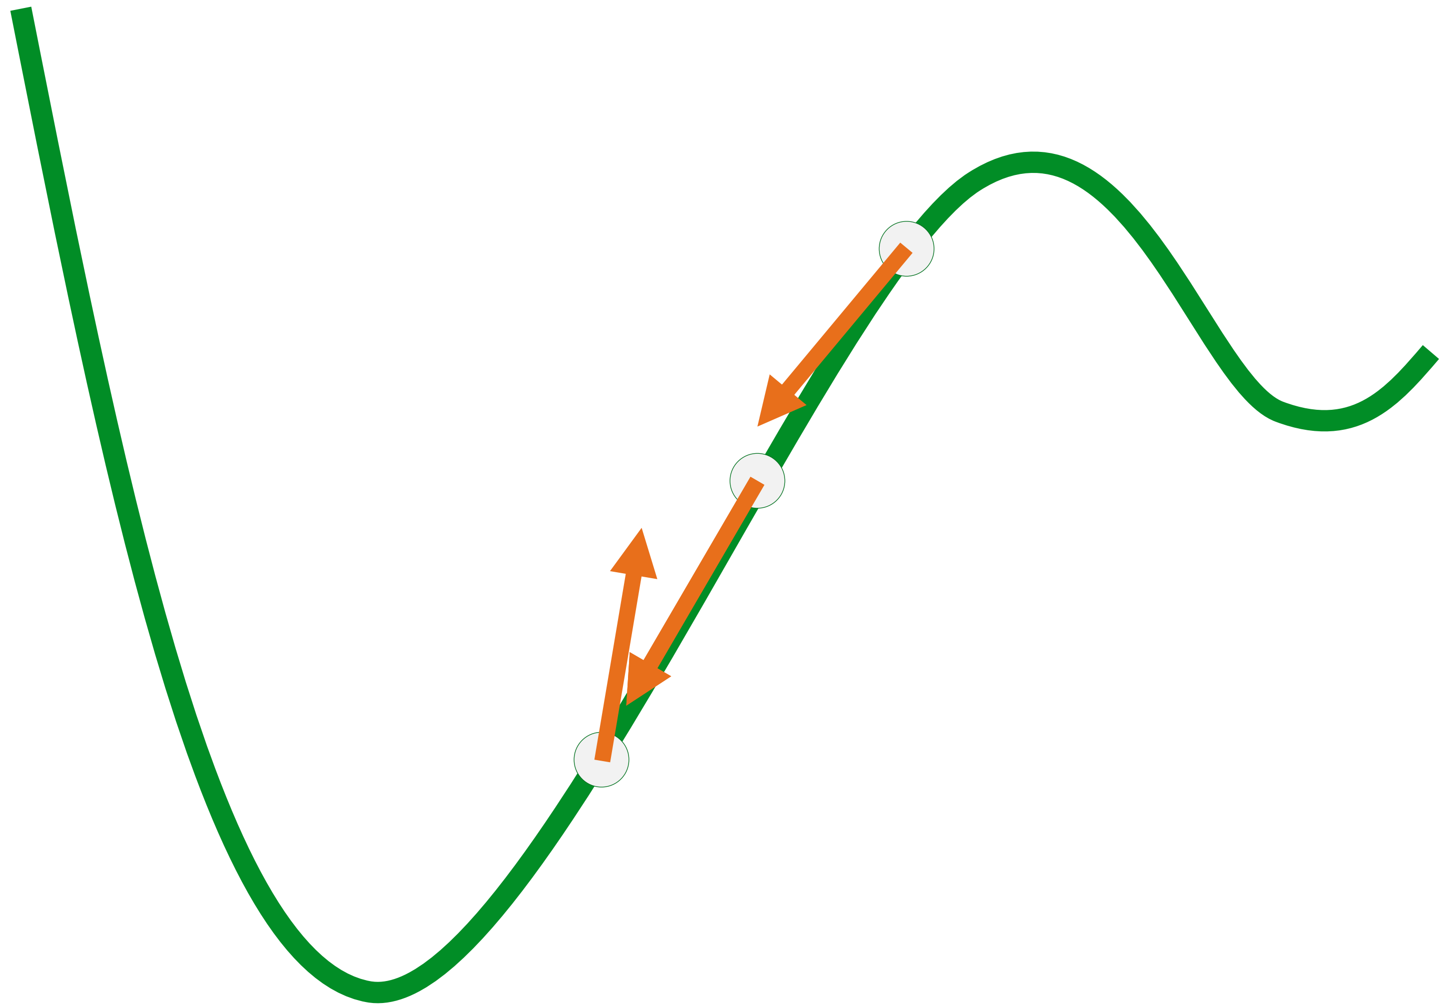

Optimizers (like Adam)

- Can smooth out stochastic gradients
- Black box

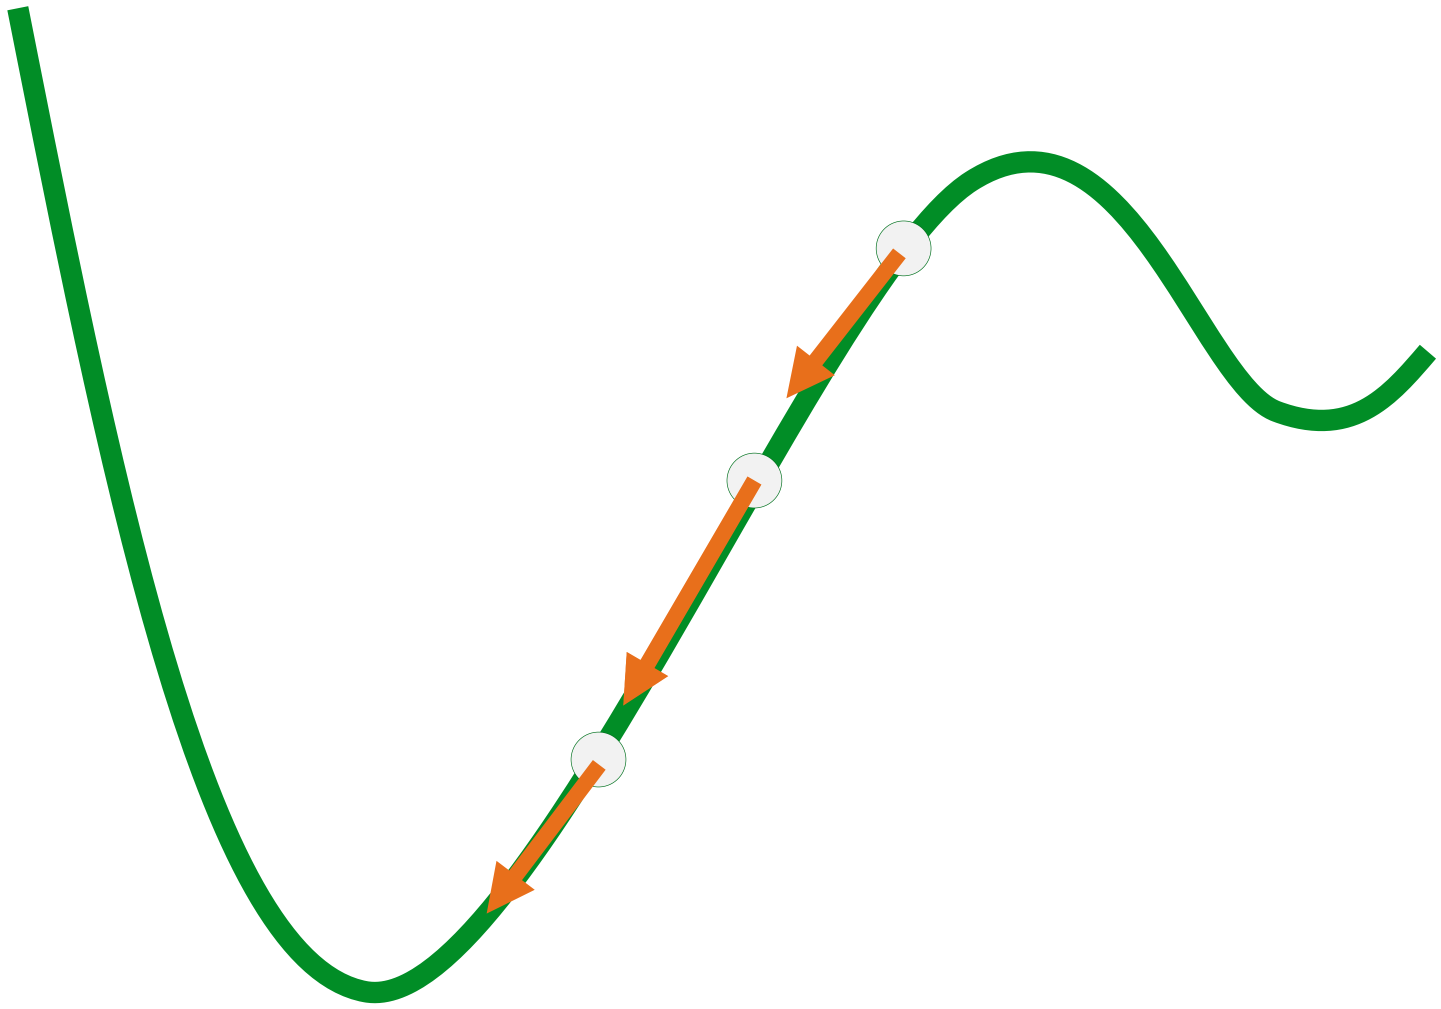

Noisy gradients

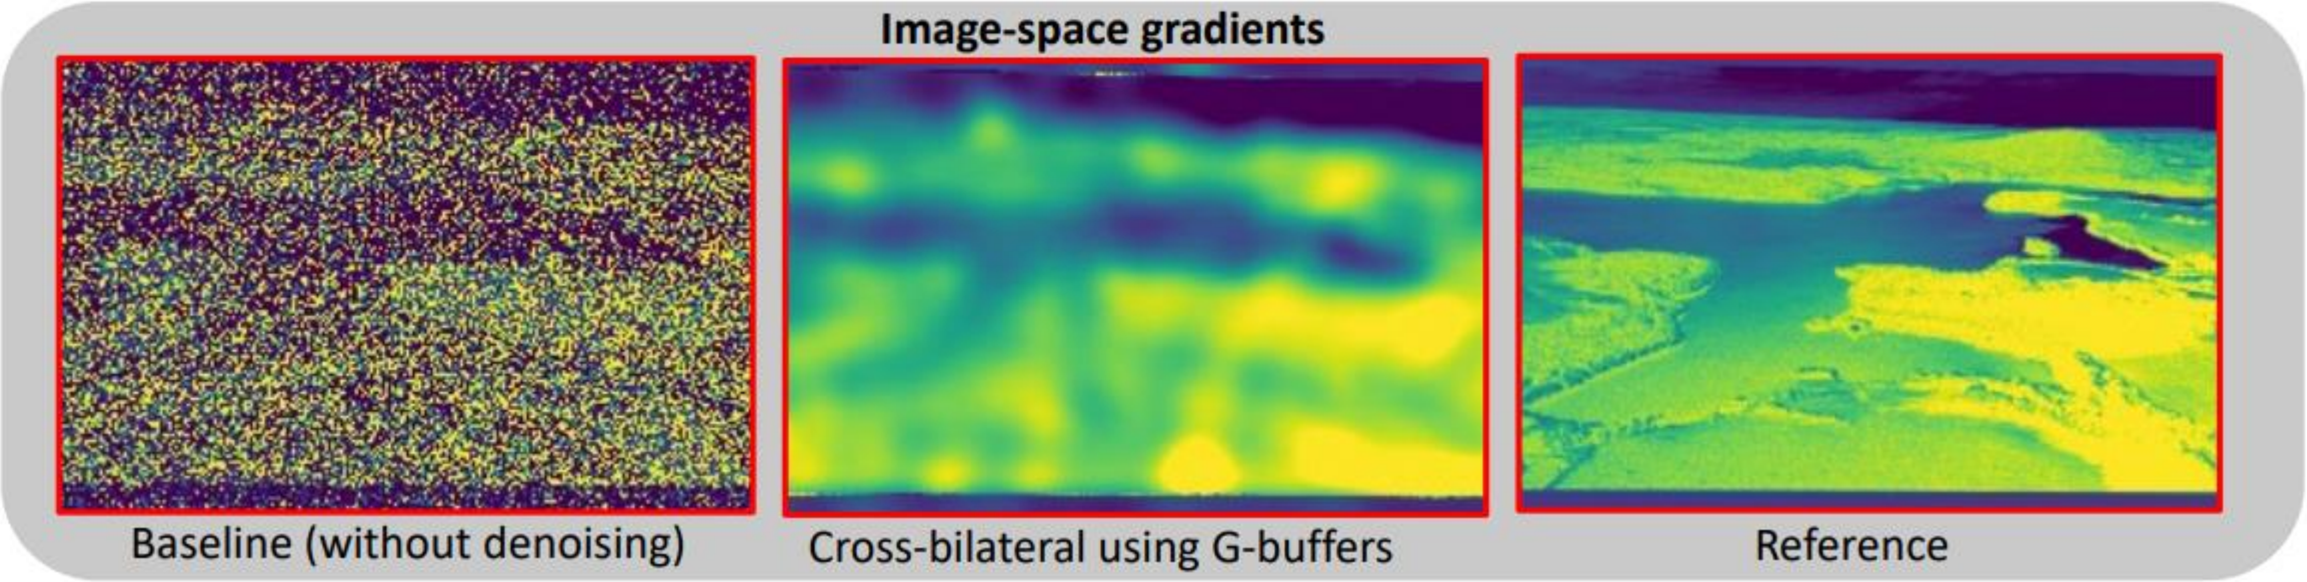

Target-Aware Image Denoising for Inverse Monte Carlo Rendering, Bochang Moon

- Spatial (target aware) filter

Bitterli et al. *SIGGRAPH 2020.* Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting.

ReSTIR (unbiased spatiotemporal reuse)

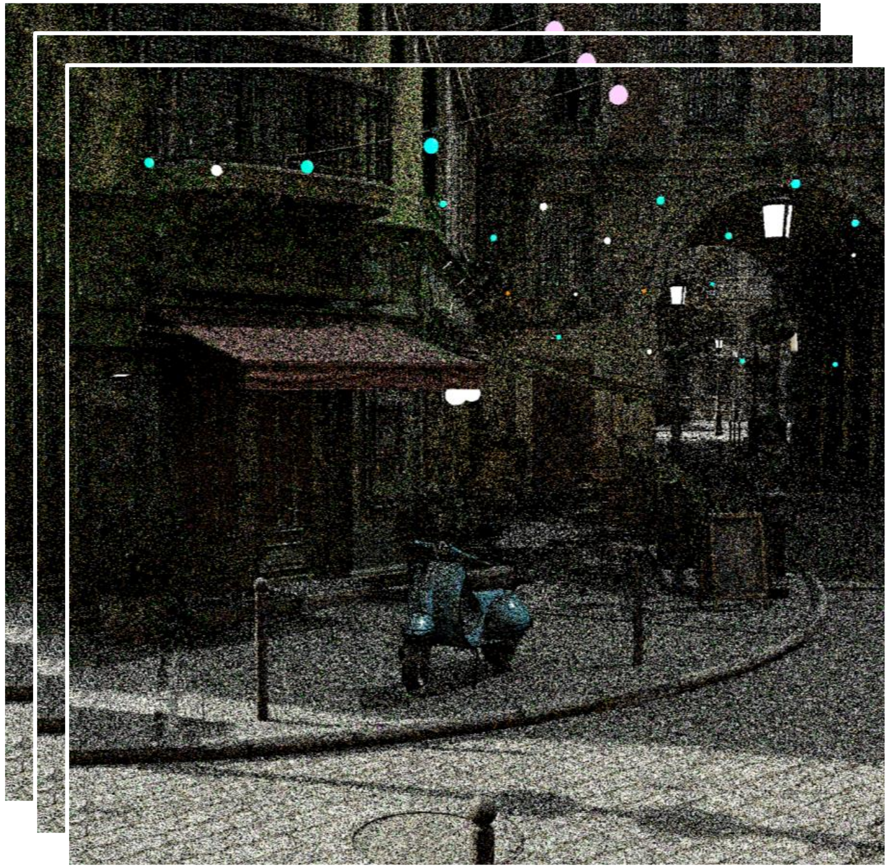

Sequence of
similar noisy frames

Reuse of
previous frames

Temporal changes are small

Initial

Target

Path Guiding – 1D example

- MC integration for $\int_0^{\pi} \sin^2 x \, dx$
  - Optimal pdf: $\frac{2}{\pi} \sin^2 x \propto \sin^2 x$

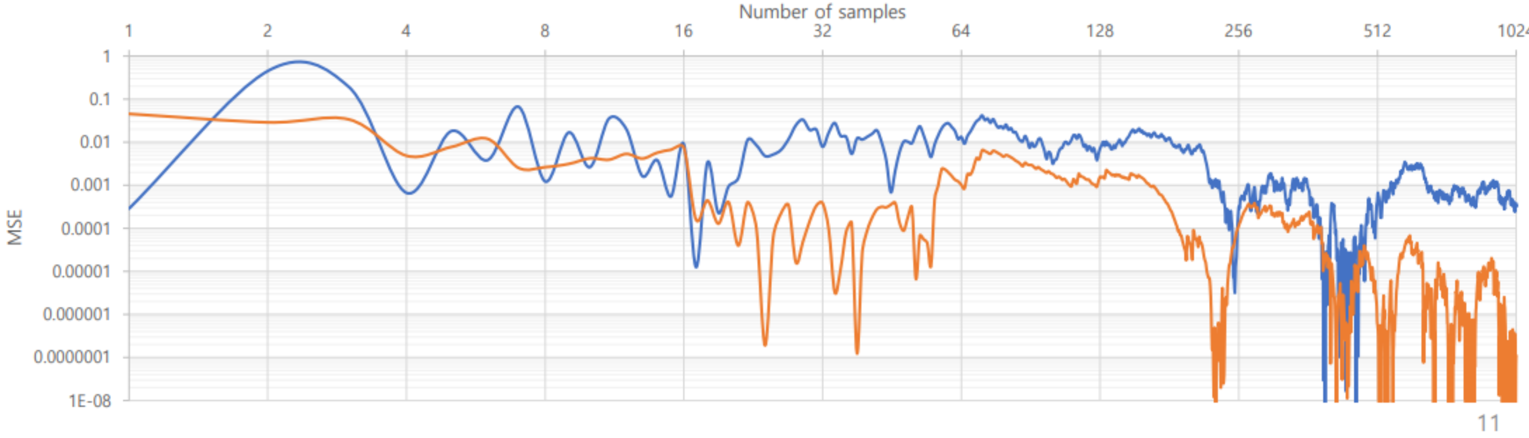

- If $p \propto L_i f_s \cos\theta_i$, $Var[\langle L_o(\mathbf{x}, \boldsymbol{\omega}_o)\rangle] = 0$

Talbot et al. *EGSR 2005*. Importance Resampling for Global Illumination.

Resampled Importance Sampling (RIS)

- Generate points Xi ~ q
- Pick one with probability proportional to f(Xi)

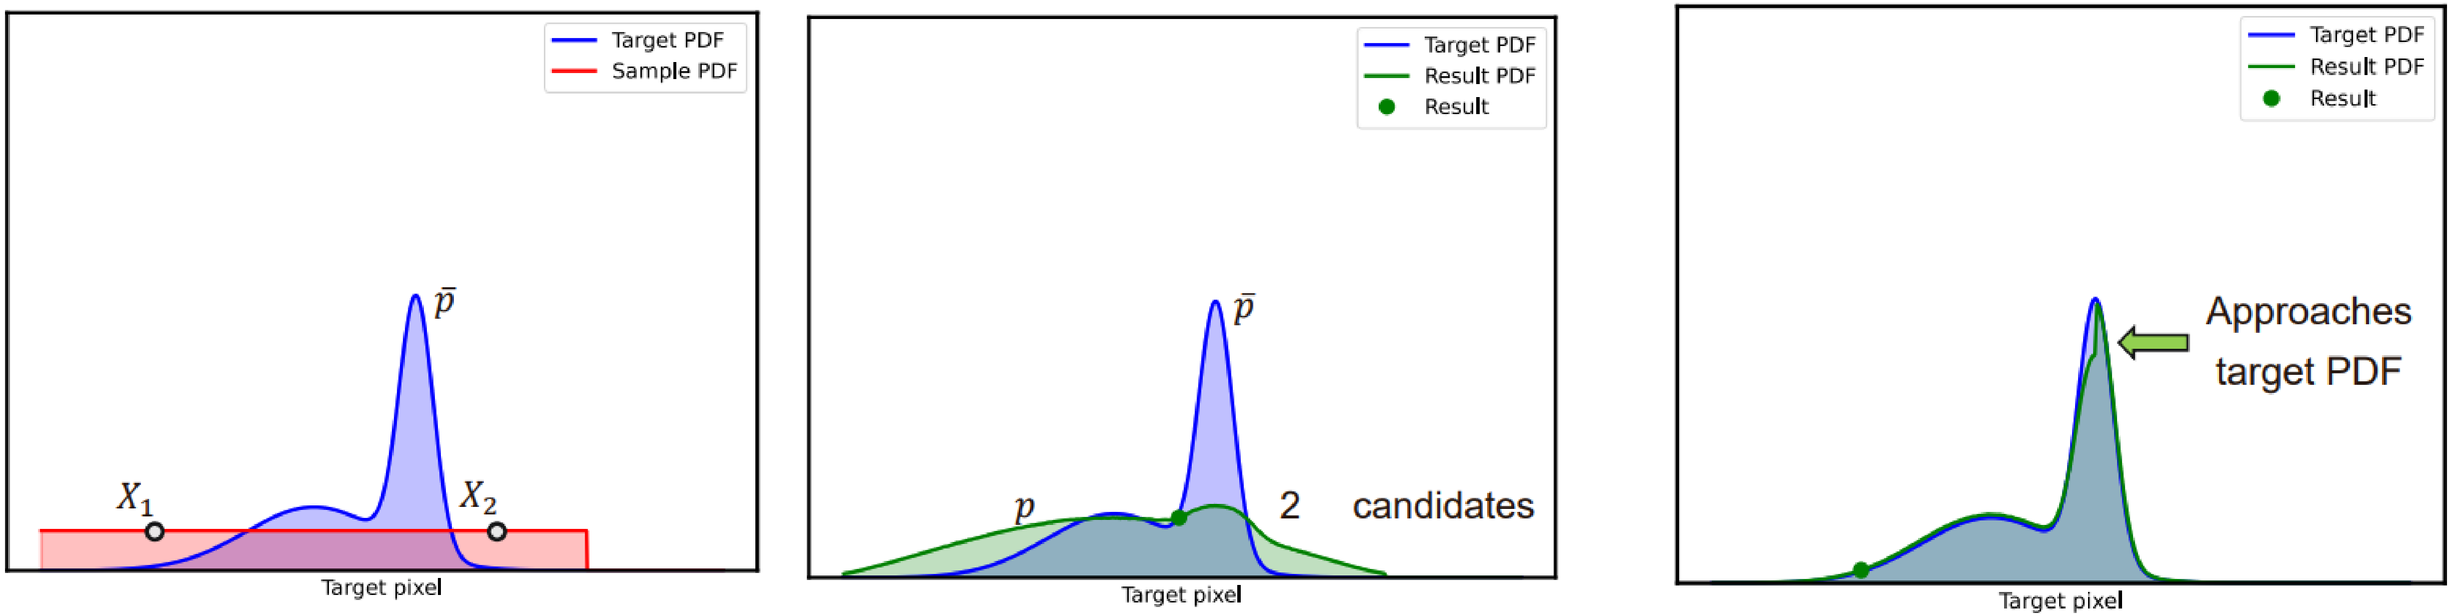

ReSTIR temporal reuse

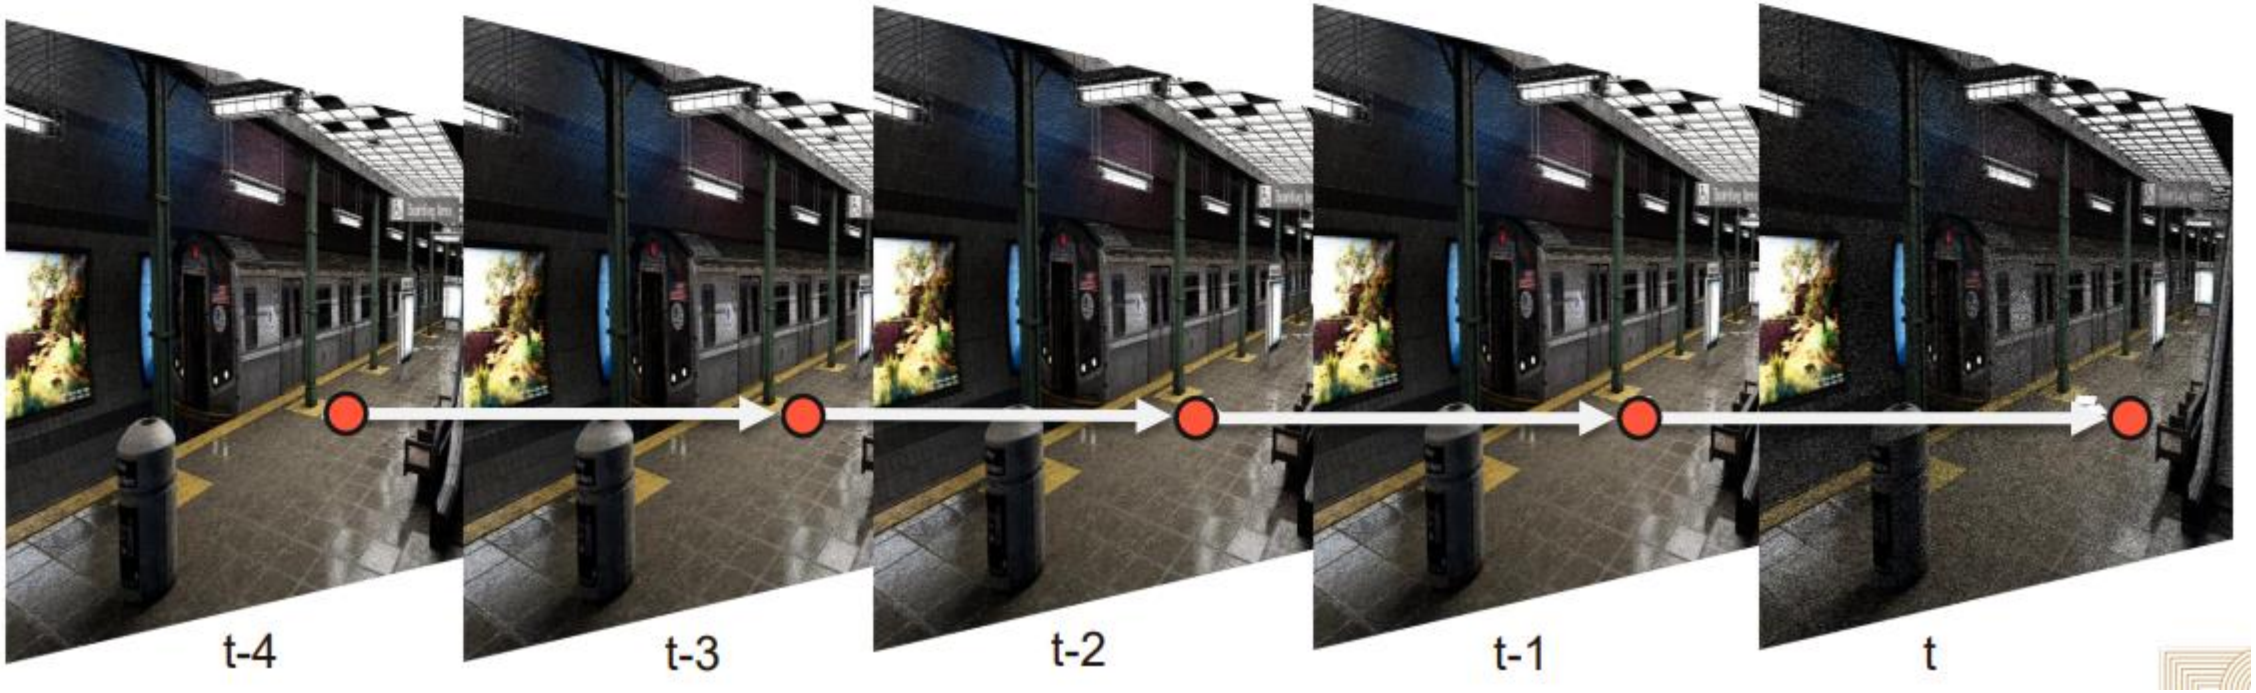

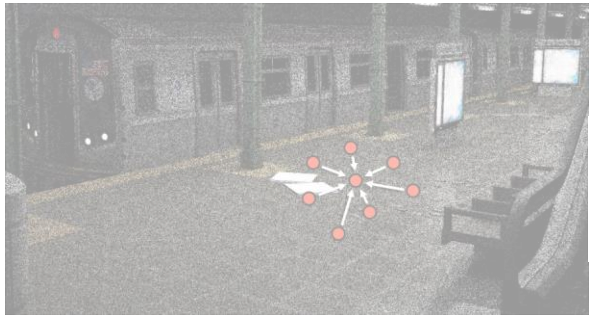

(spatial reuse is not used in ReSTIR DR)

Can we just apply ReSTIR then?

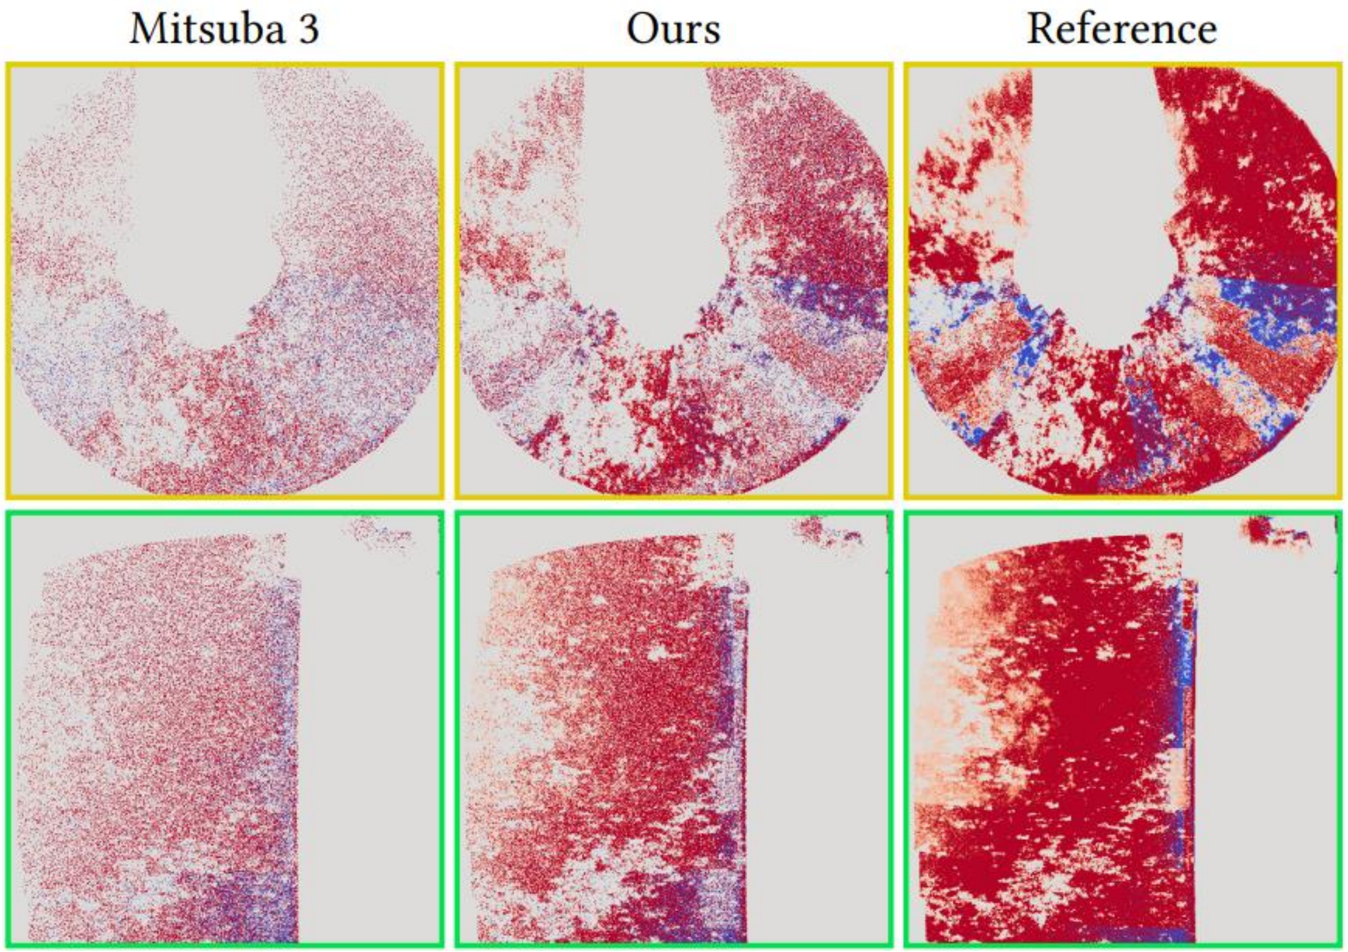

Theoretical contributions

- Parameter-Space Differentiable Rendering
- Resampling with Positive and Negative Functions
    - Positivization

The Problem with Pixel-centric Differentiable Rendering

Forward Rendering

Single intensity I for each of N pixels = N samples

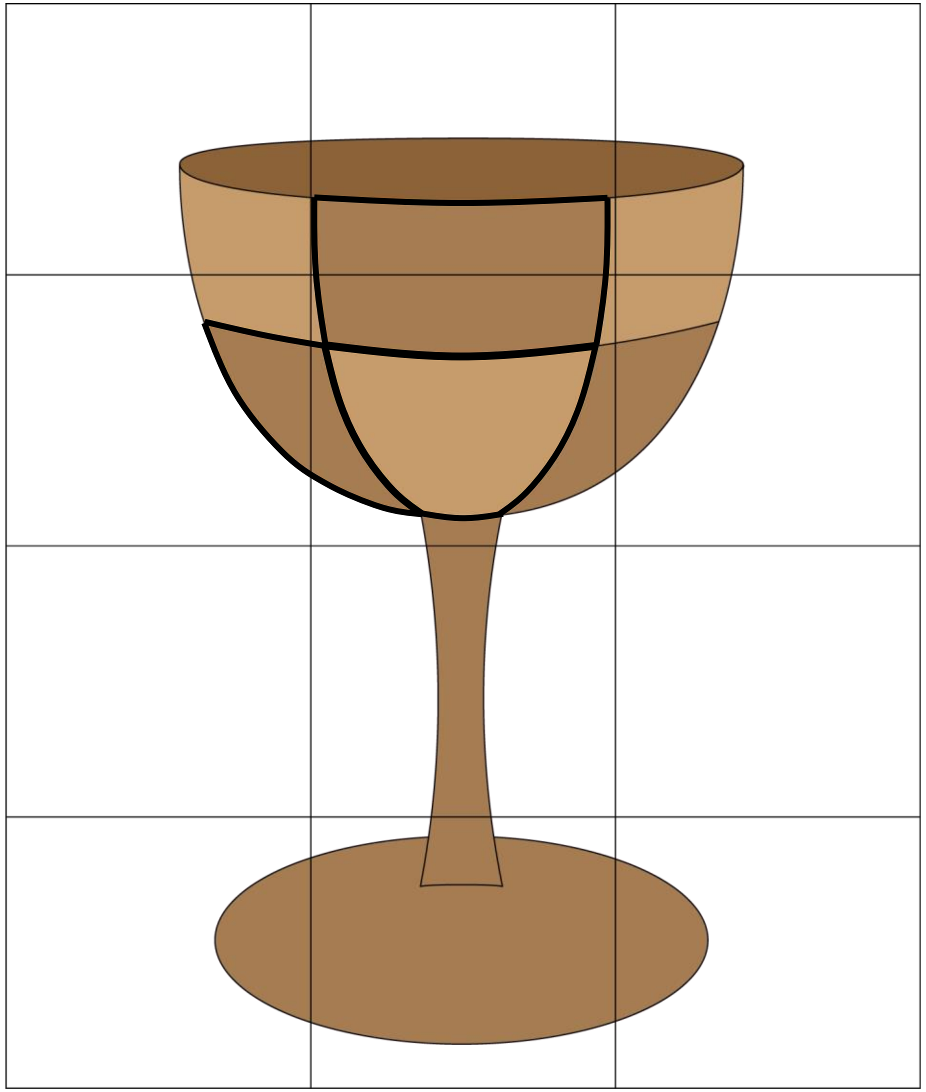

Differentiable Rendering

One derivative for each texel π_i in each pixel

$$\frac{\partial I}{\partial \pi_0} \quad \frac{\partial I}{\partial \pi_1} \quad \frac{\partial I}{\partial \pi_2} \dots$$

M texels = N · M samples

Parameter-Space ReSTIR

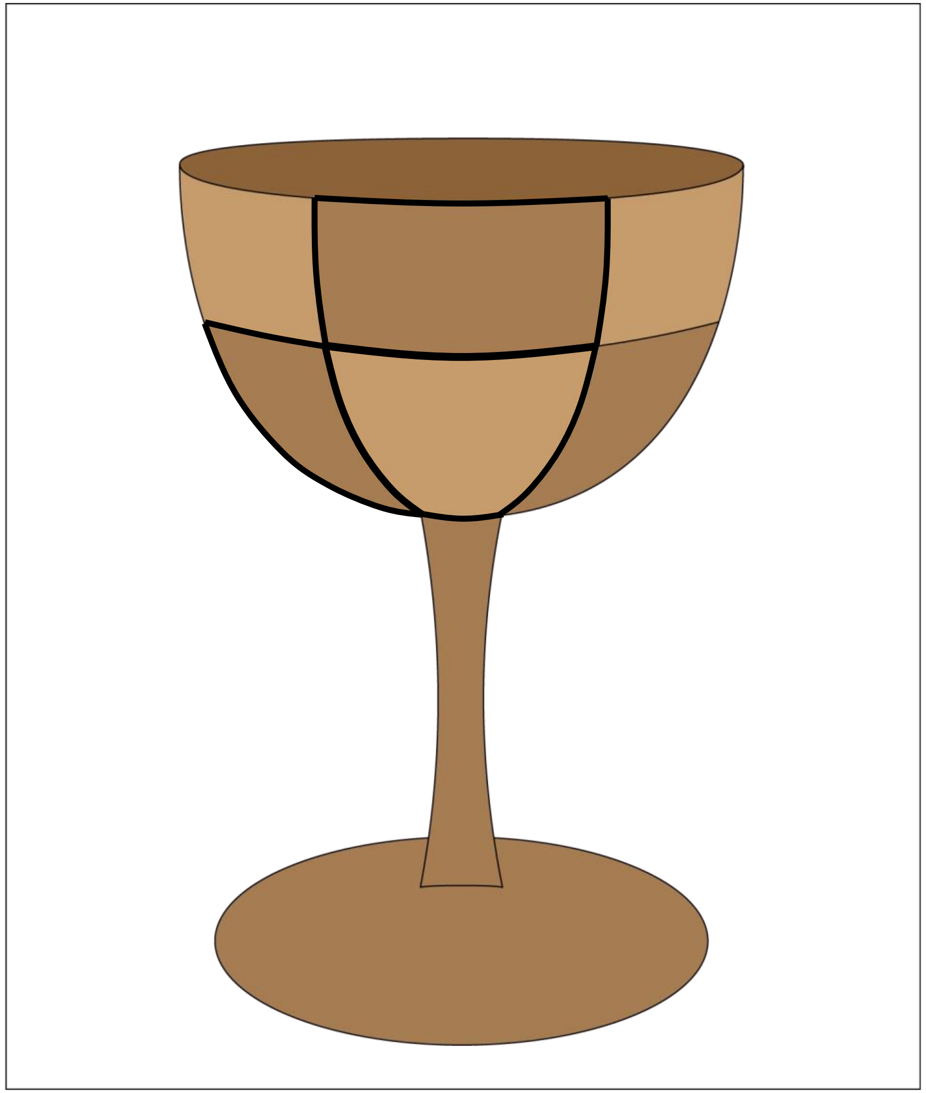

Theoretical contributions

- Parameter-Space Differentiable Rendering
- Resampling with Positive and Negative Functions
  - Positivization

Positive and negative function

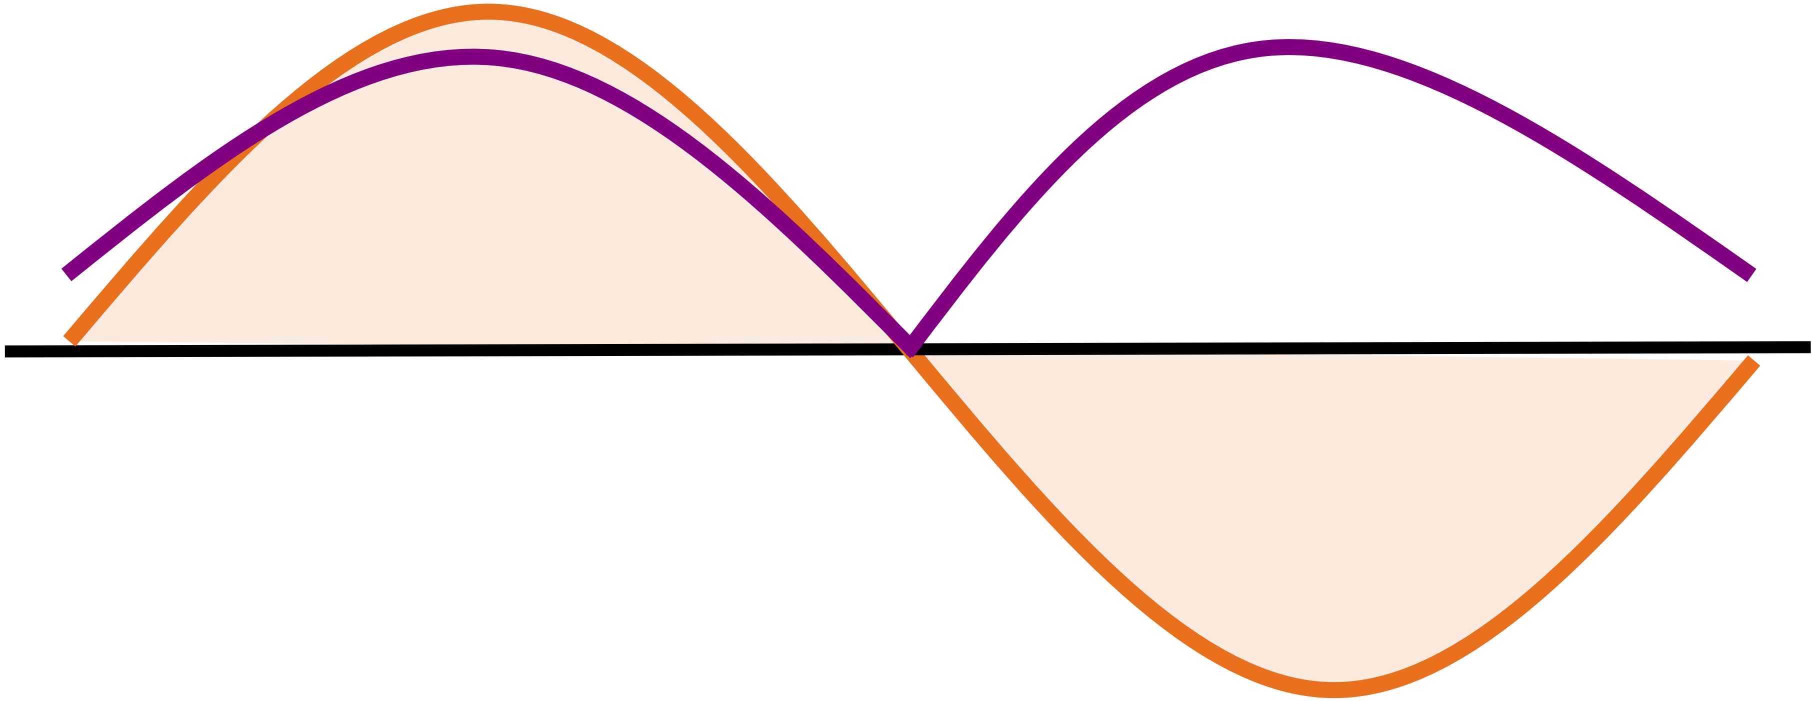

$$V\left[\frac{\partial f}{|\partial f|}\right] \neq 0$$

Noise due to sign

Owen et al. *Journal of the American Statistical Association (2000)*
Safe and effective importance sampling.

Positivization

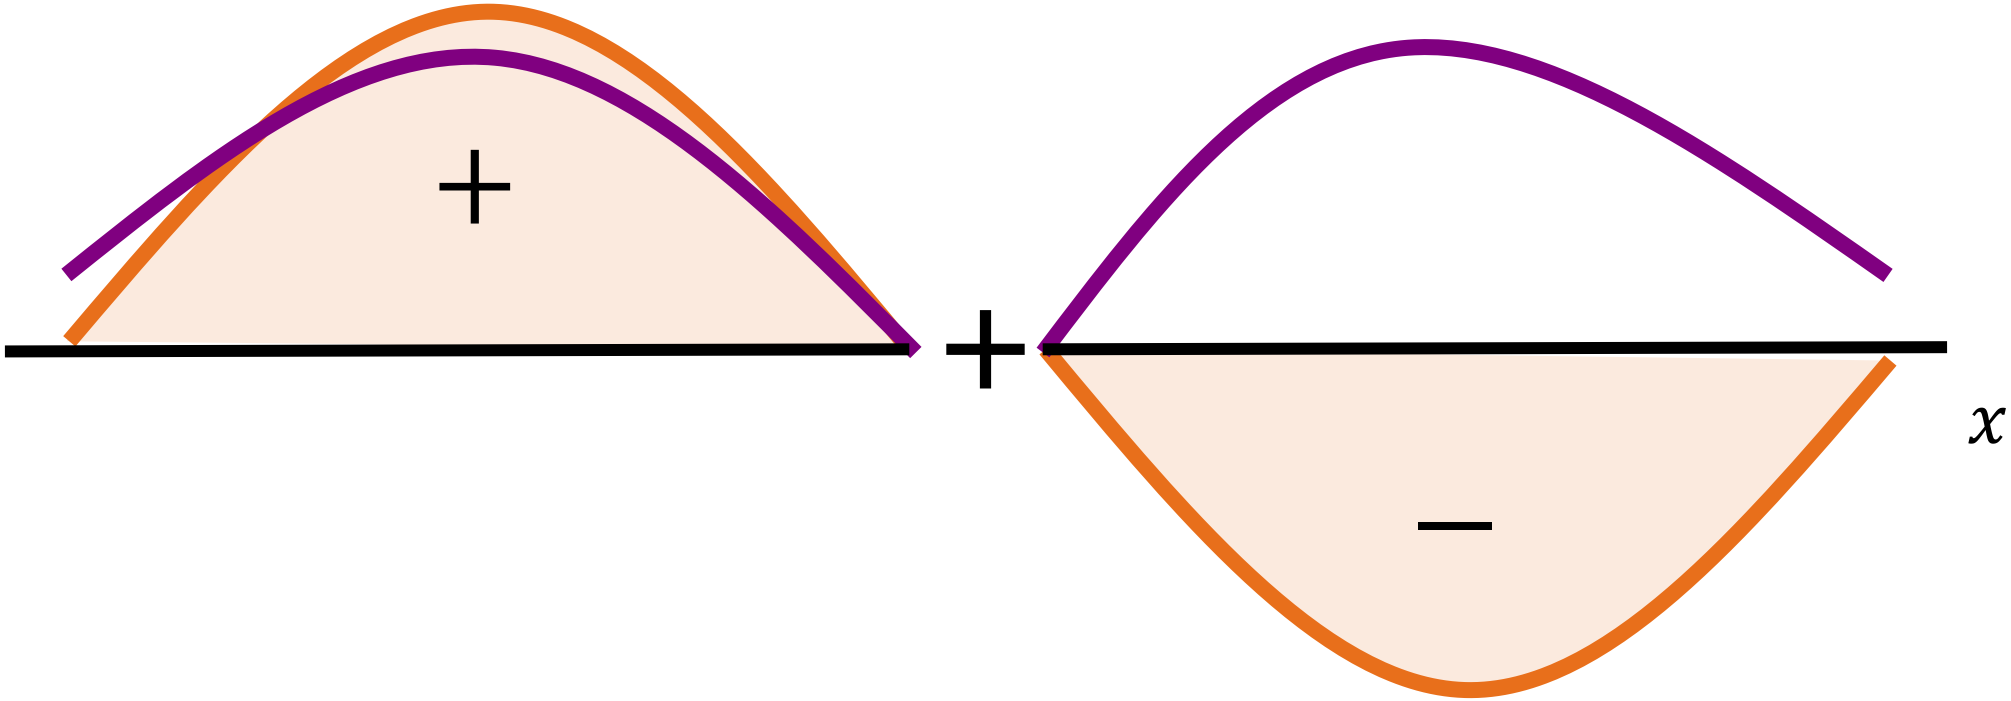

Sample q_+, q_-

Variance → 0 when

$$q_+ = \max(\partial f, 0) \quad q_- = \max(-\partial f, 0)$$

Theoretical contributions

- Parameter-Space Differentiable Rendering
- Resampling with Positive and Negative Functions
  - Positivization

Our Texture Optimization Algorithm

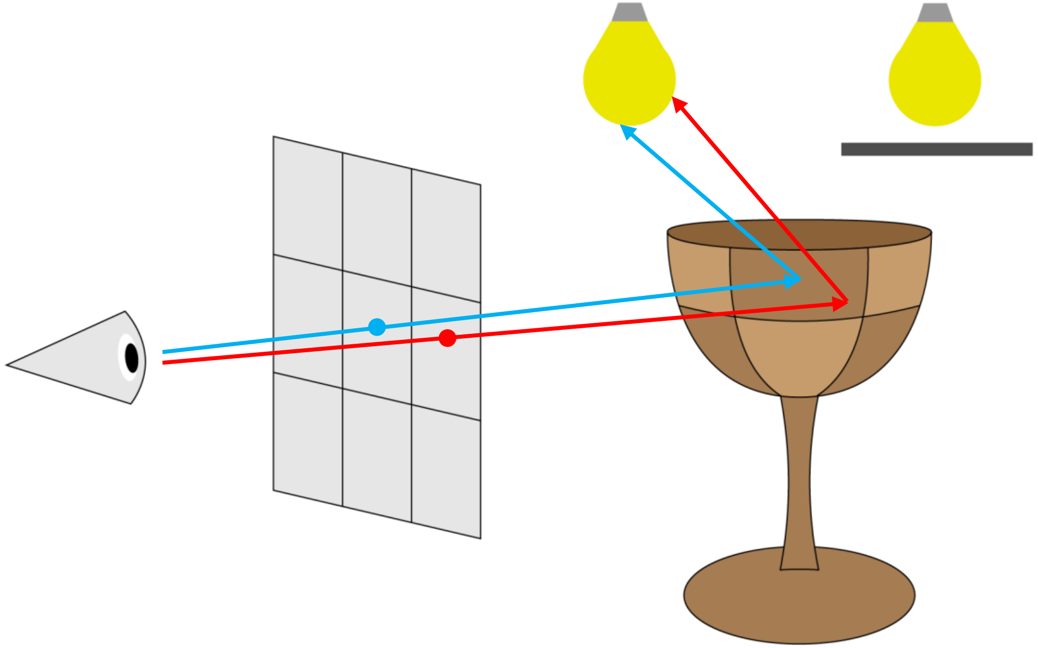

Store
1 positive, 1 negative
sample per texel

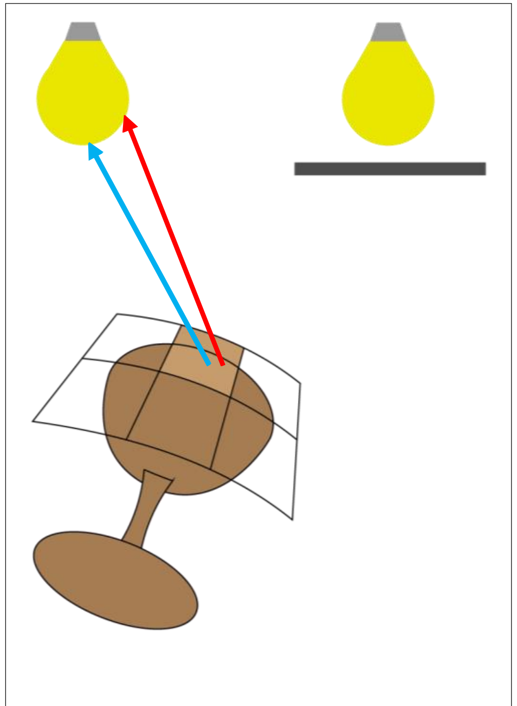

Iteration i

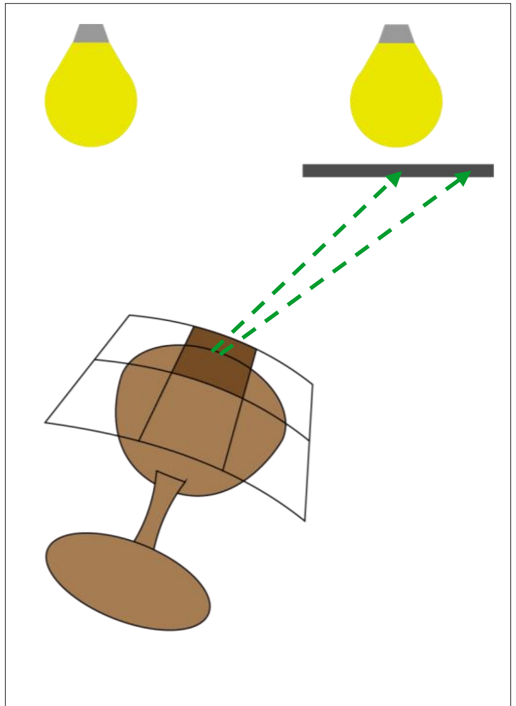

Iteration $i + 1$

Result

implemented method on top of a direct lighting integrator in Mitsuba 3

the recovery of the roughness texture of the chalice with many colored lights

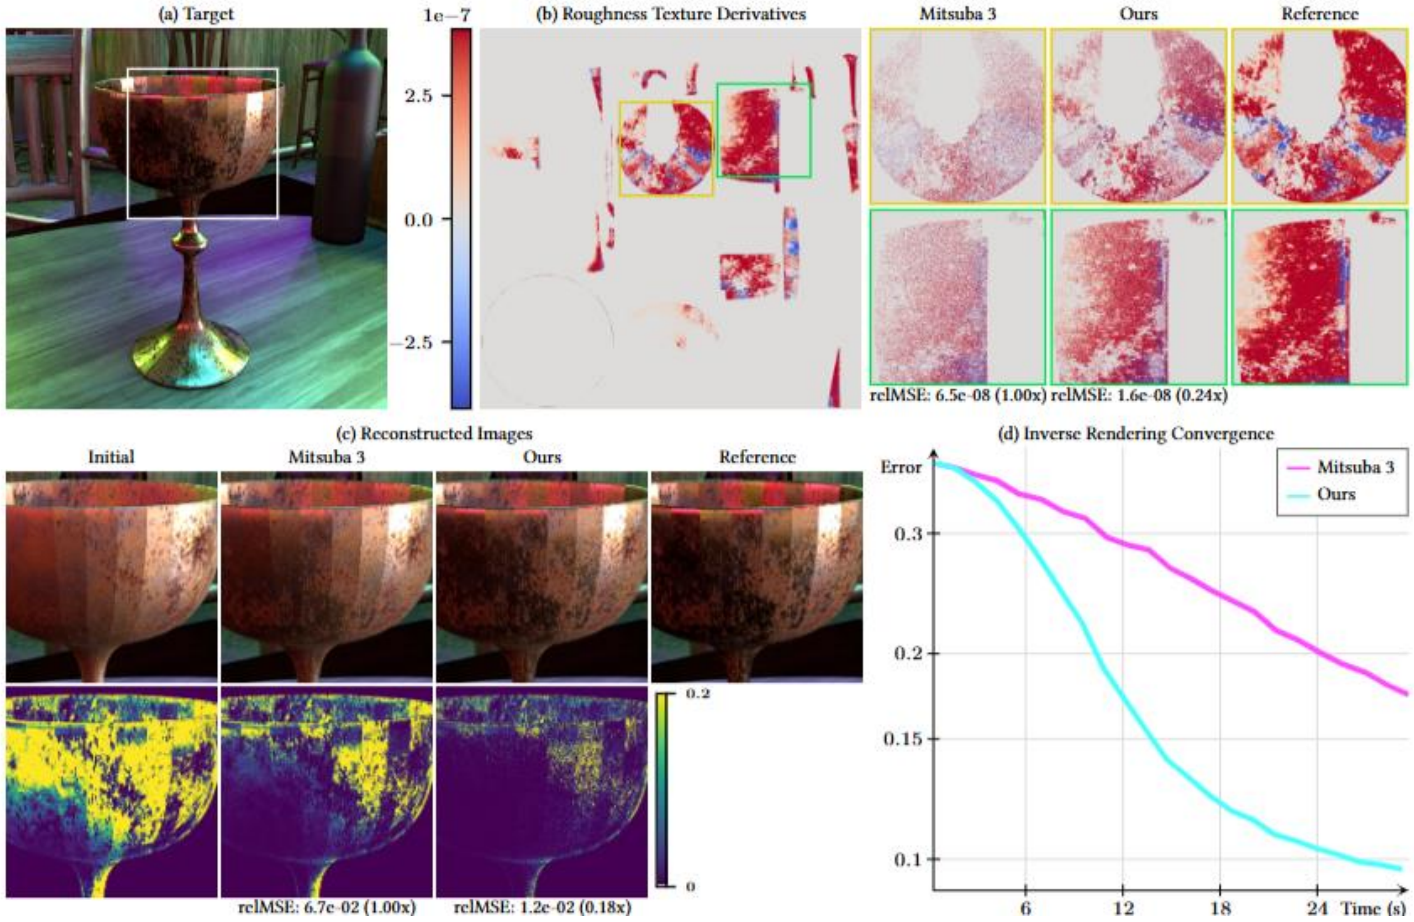

All experiments ran on an NVIDIA GeForce RTX 2080 Ti.

Results: Gradients – Disney BSDF Roughness

Initial

Target

Texture Gradient

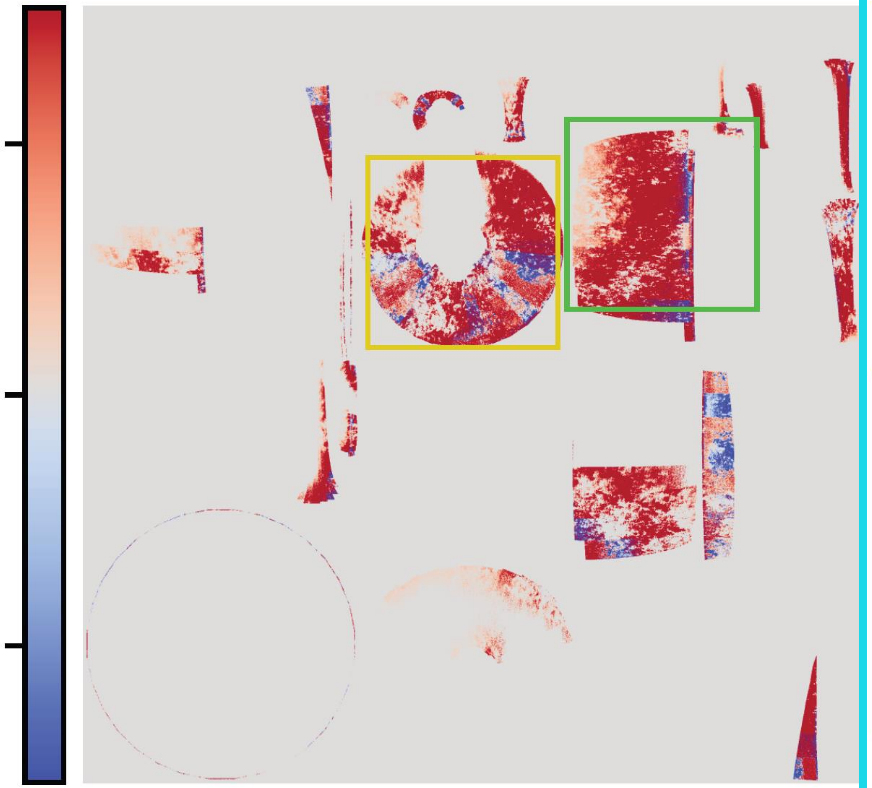

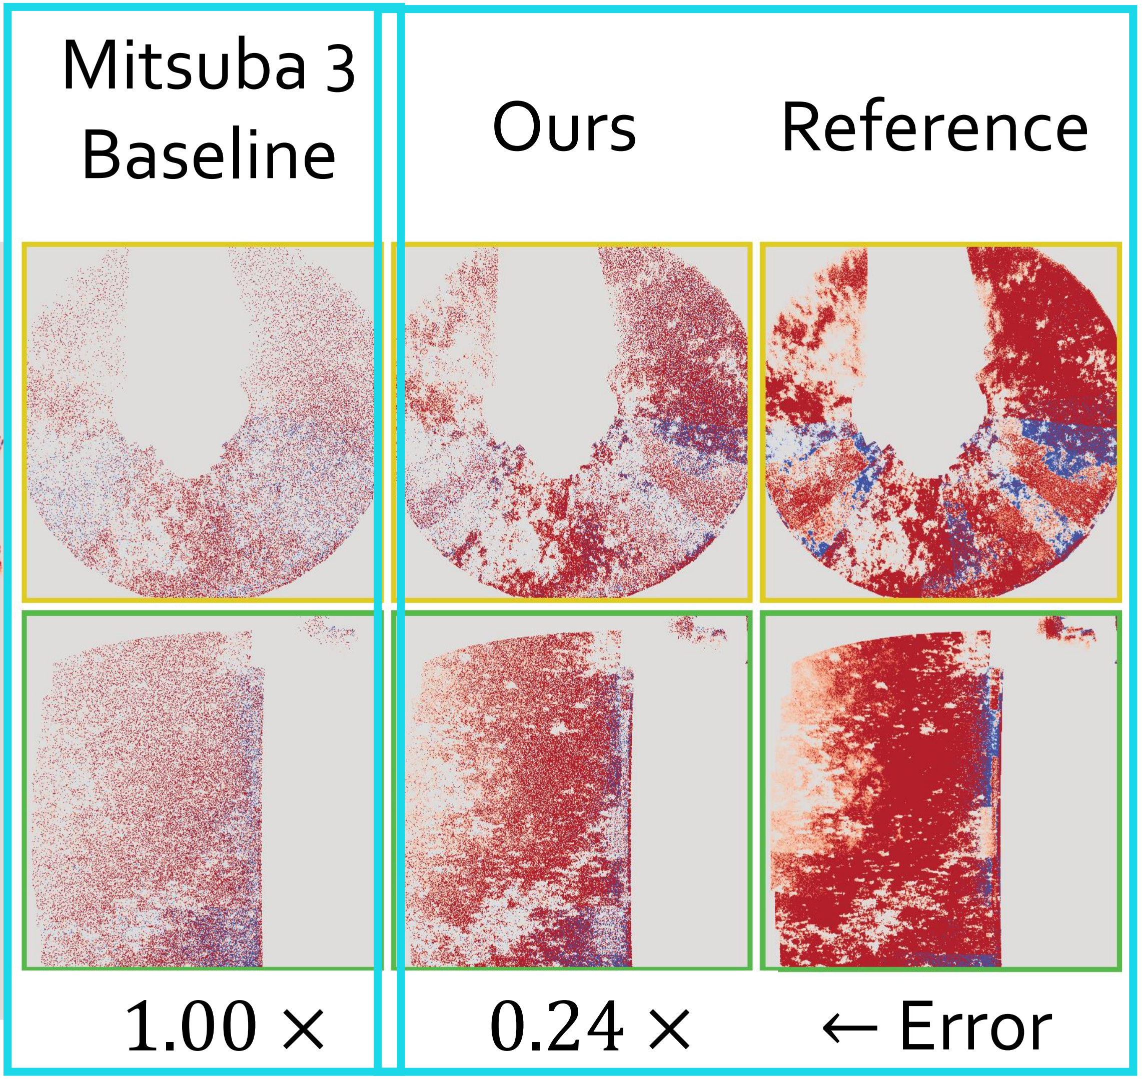

high variance computes
noisy gradients

Gradients – Disney BSDF Roughness

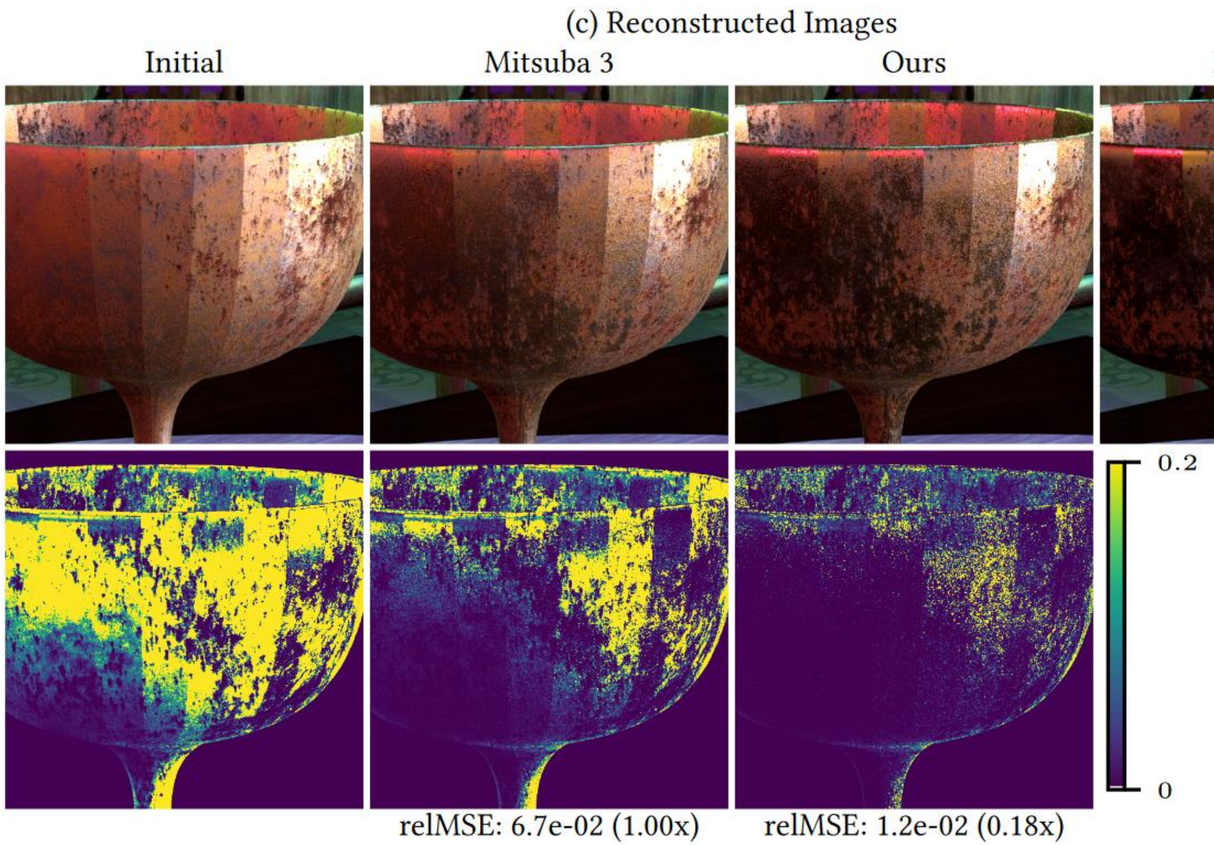

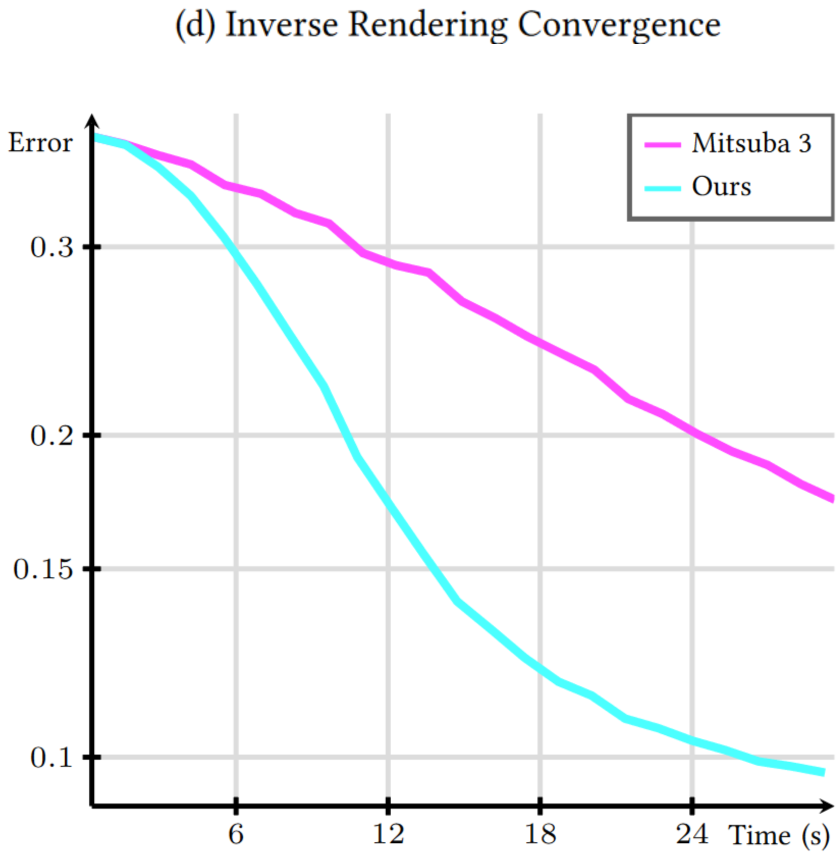

reduces variance of gradient estimates, resulting **faster** inverse rendering

Results: Inverse Rendering – Disney BSDF Anisotropy

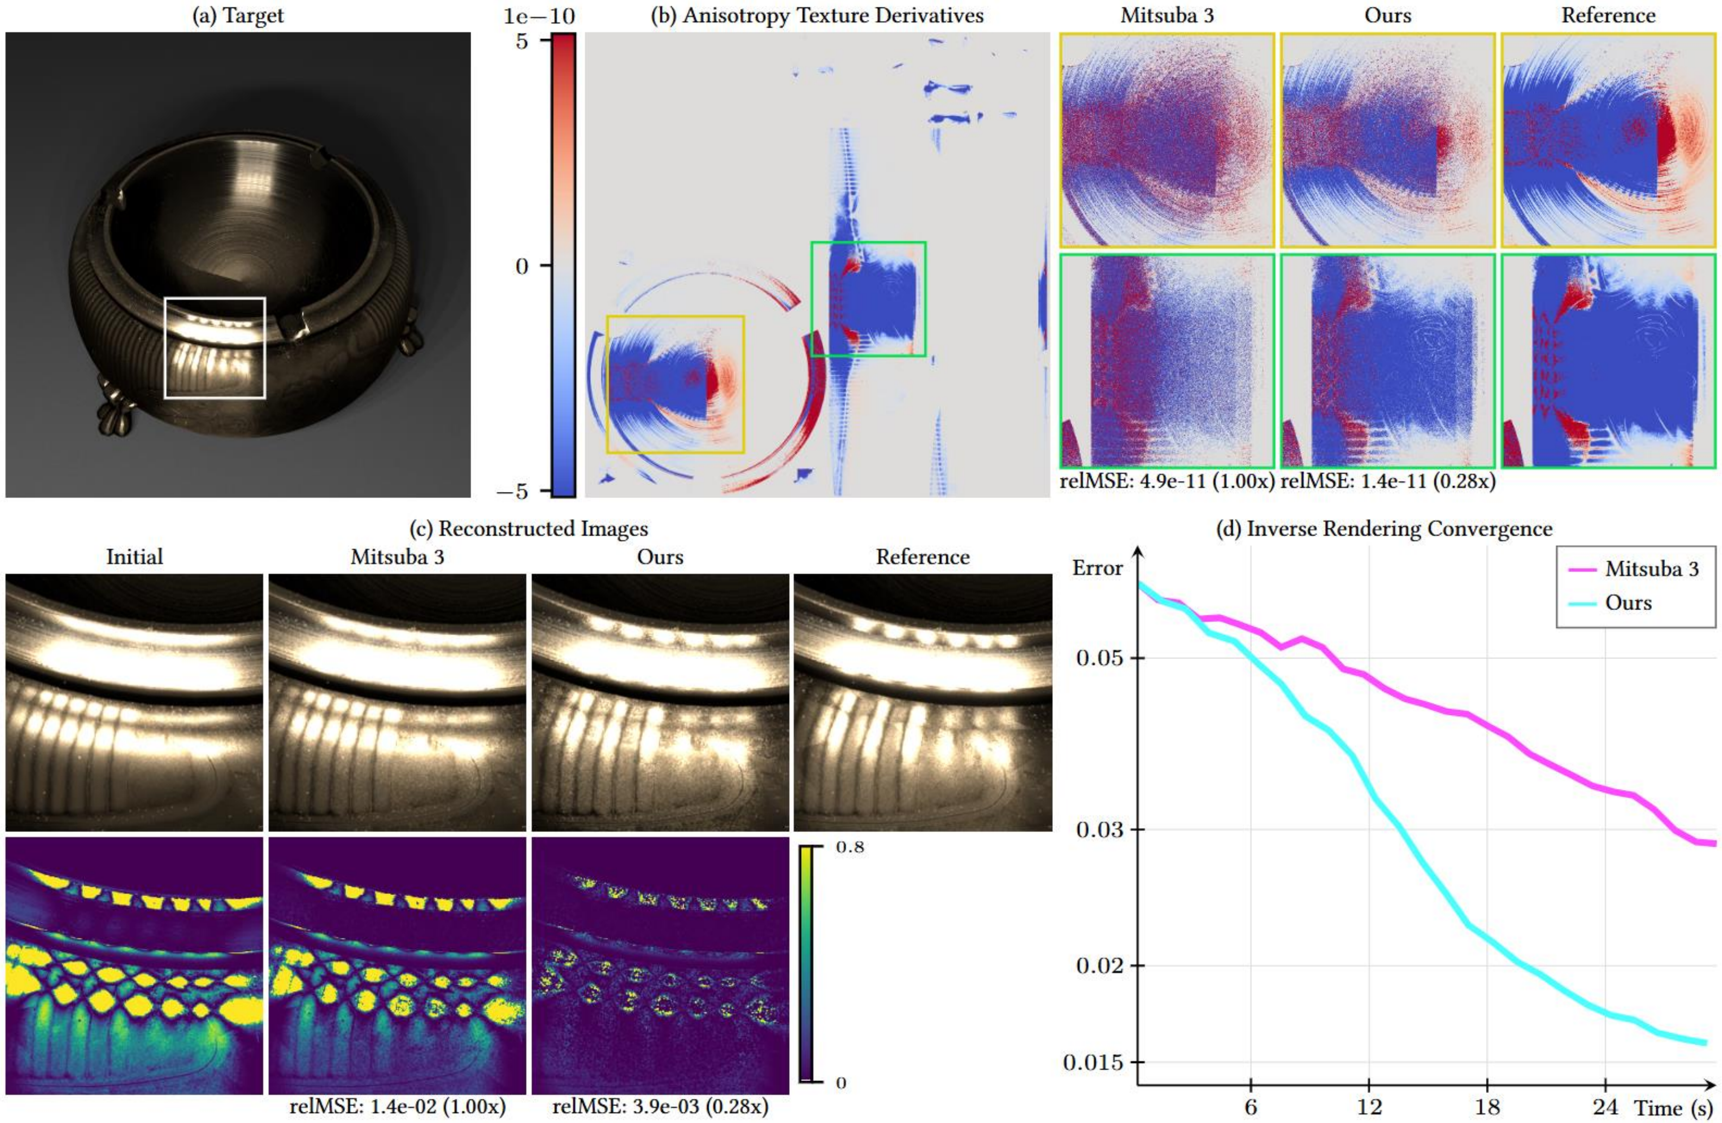

Results: Inverse Rendering – Disney BSDF Anisotropy

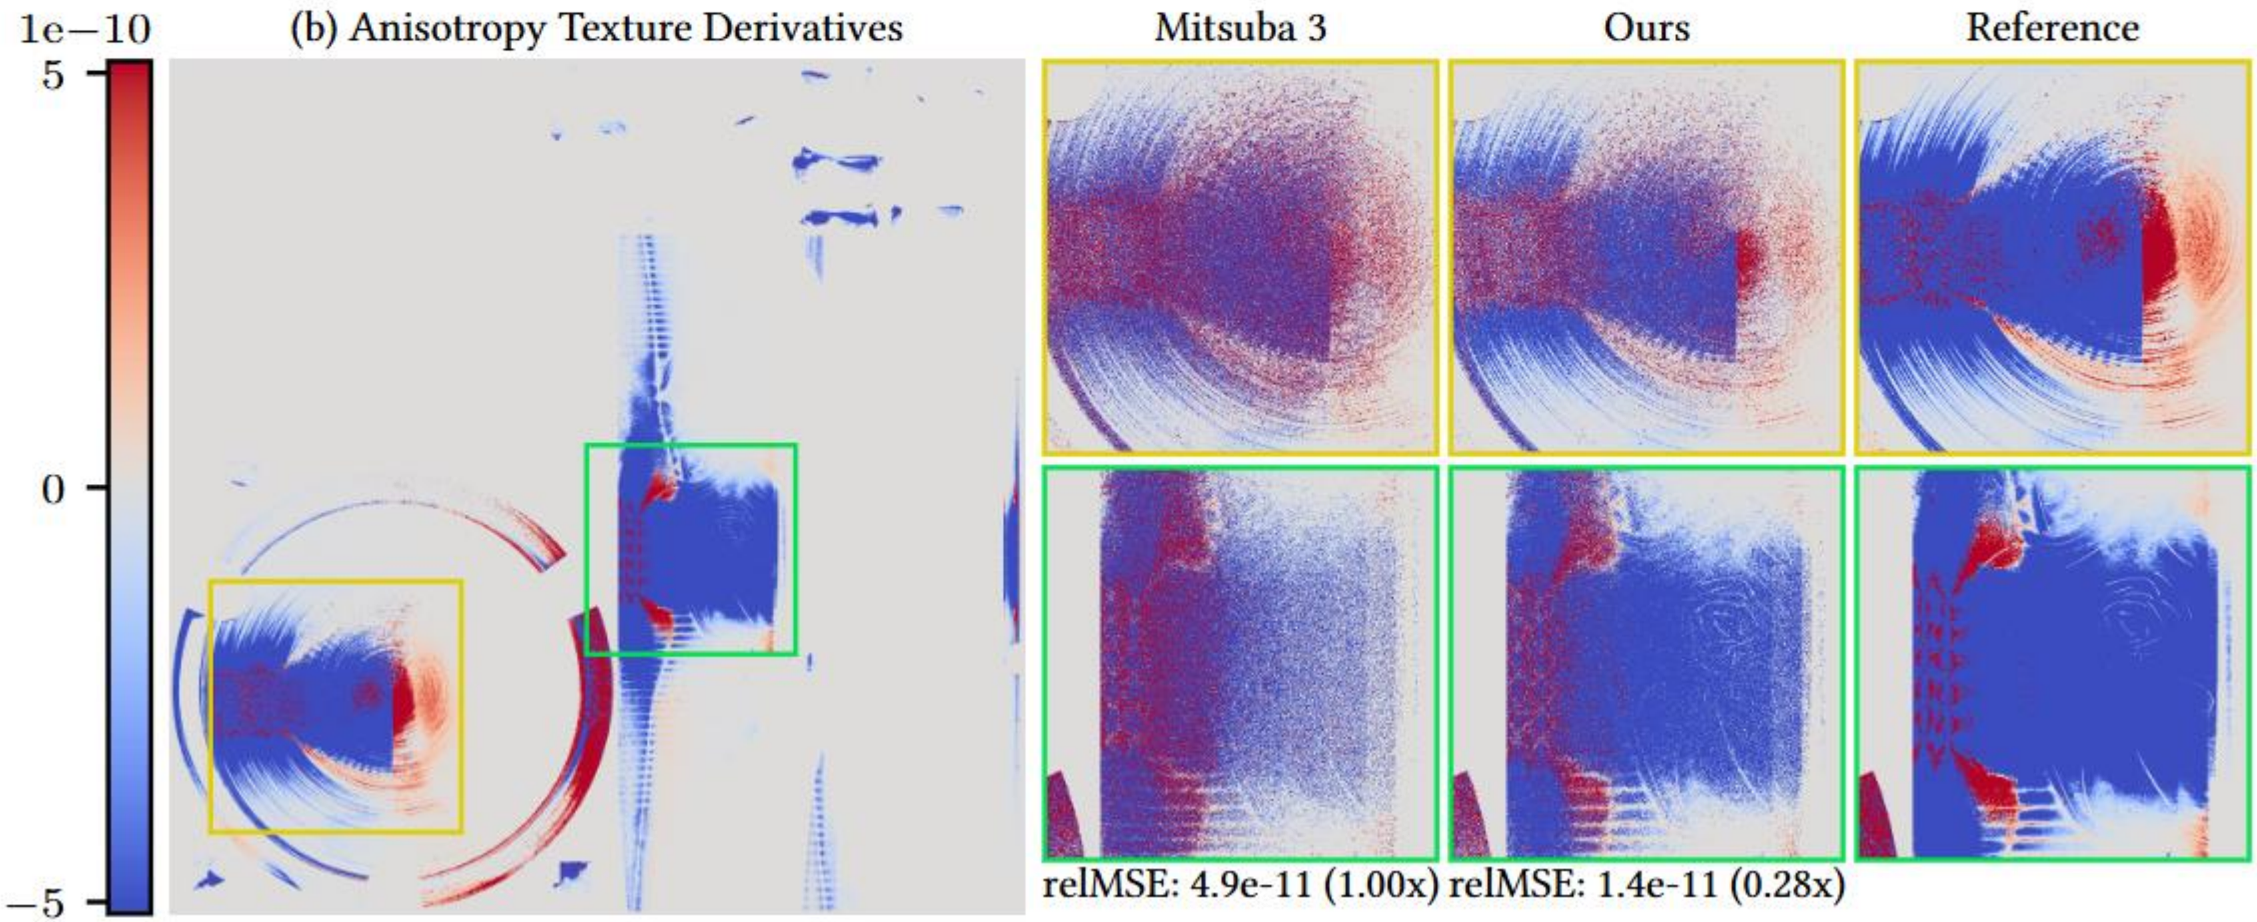

derivatives with
the wrong sign

Results: Inverse Rendering – Disney BSDF Anisotropy

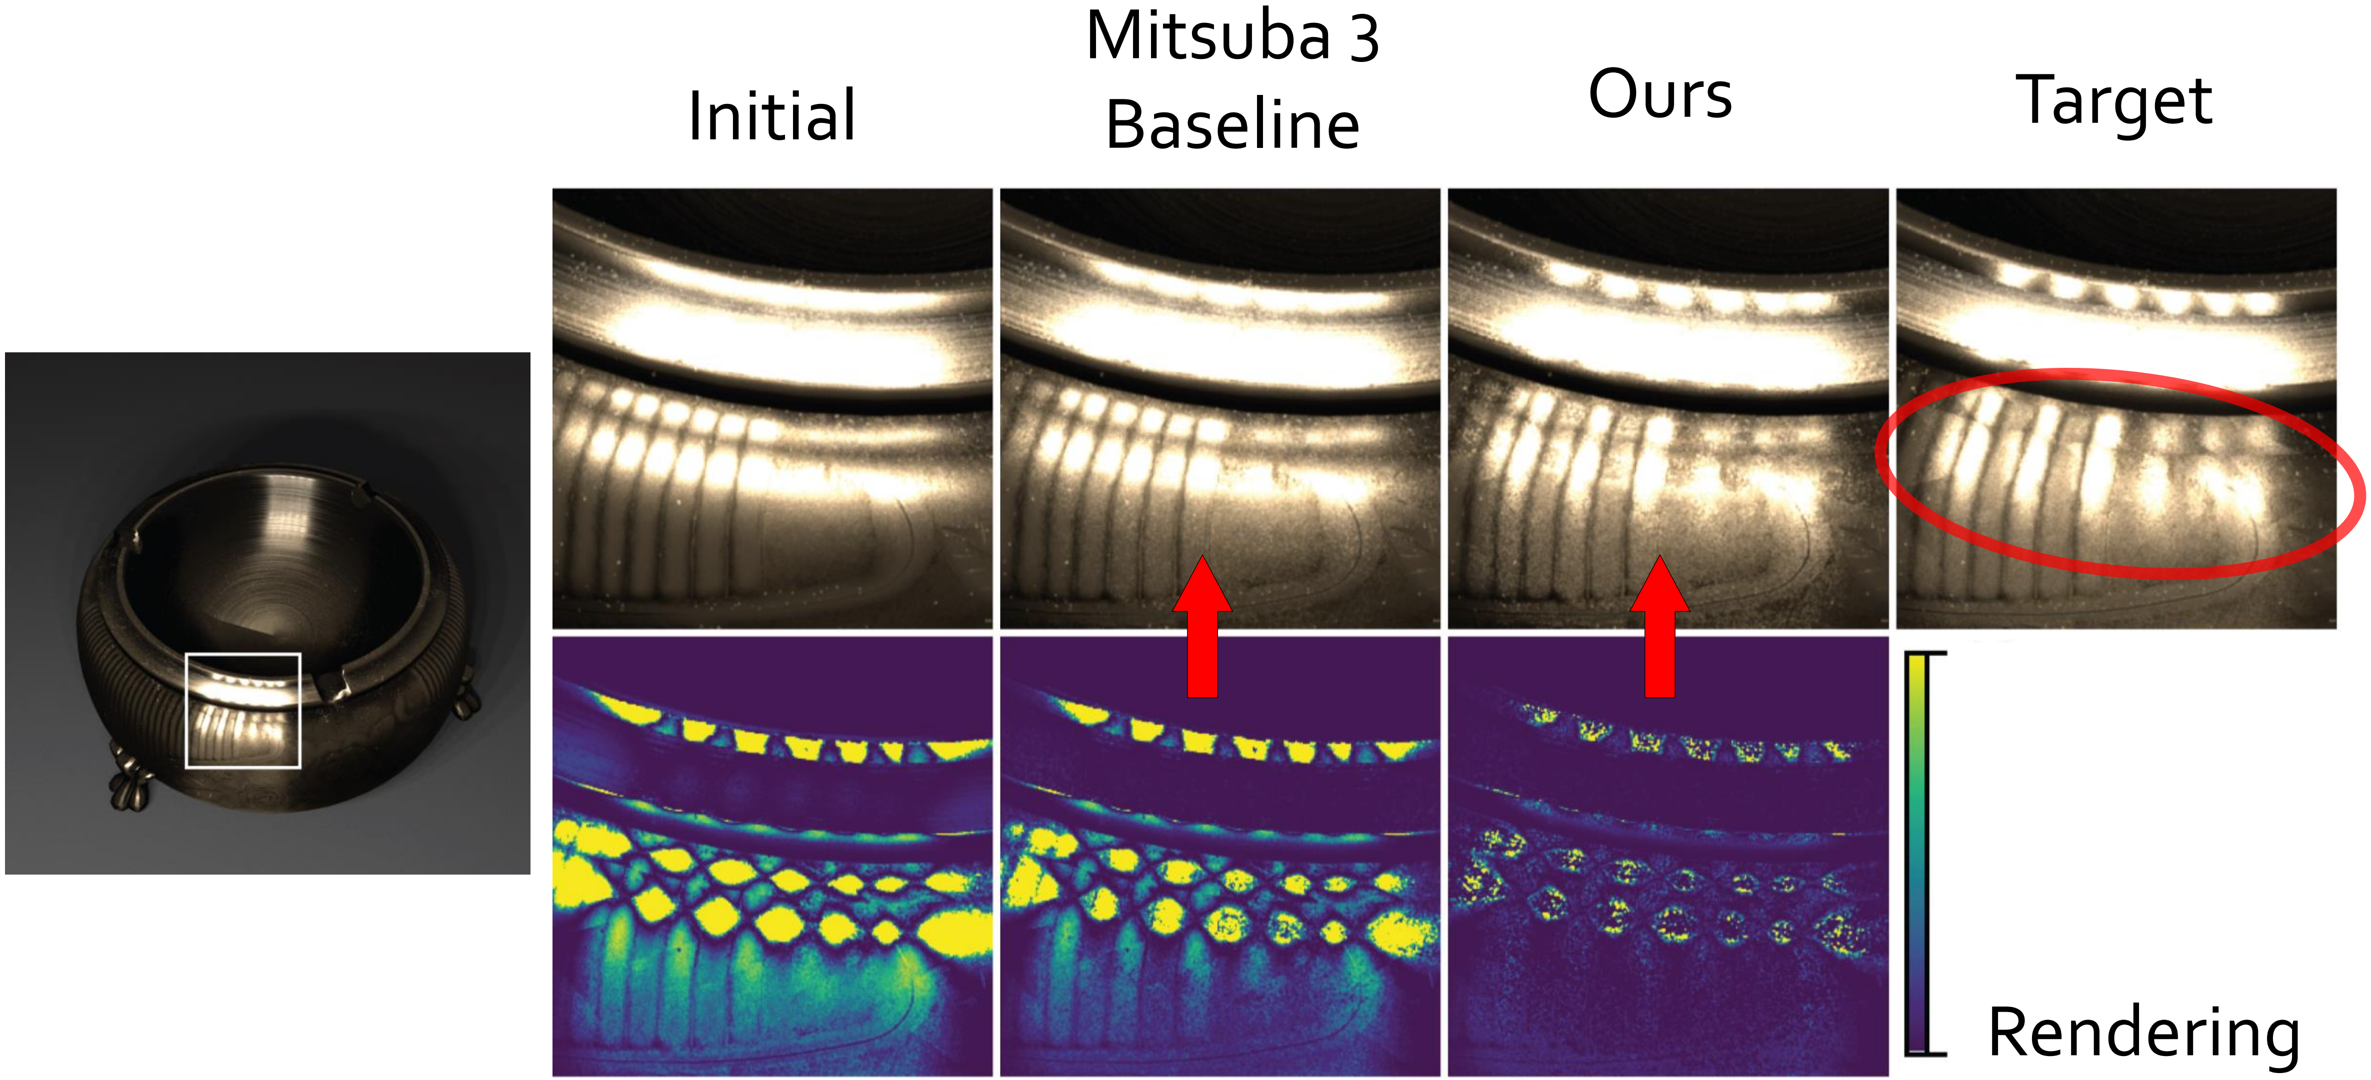

$2-3\times$ Faster Convergence $1.00\times$ $0.28\times$

Results: Positivized (G)RIS

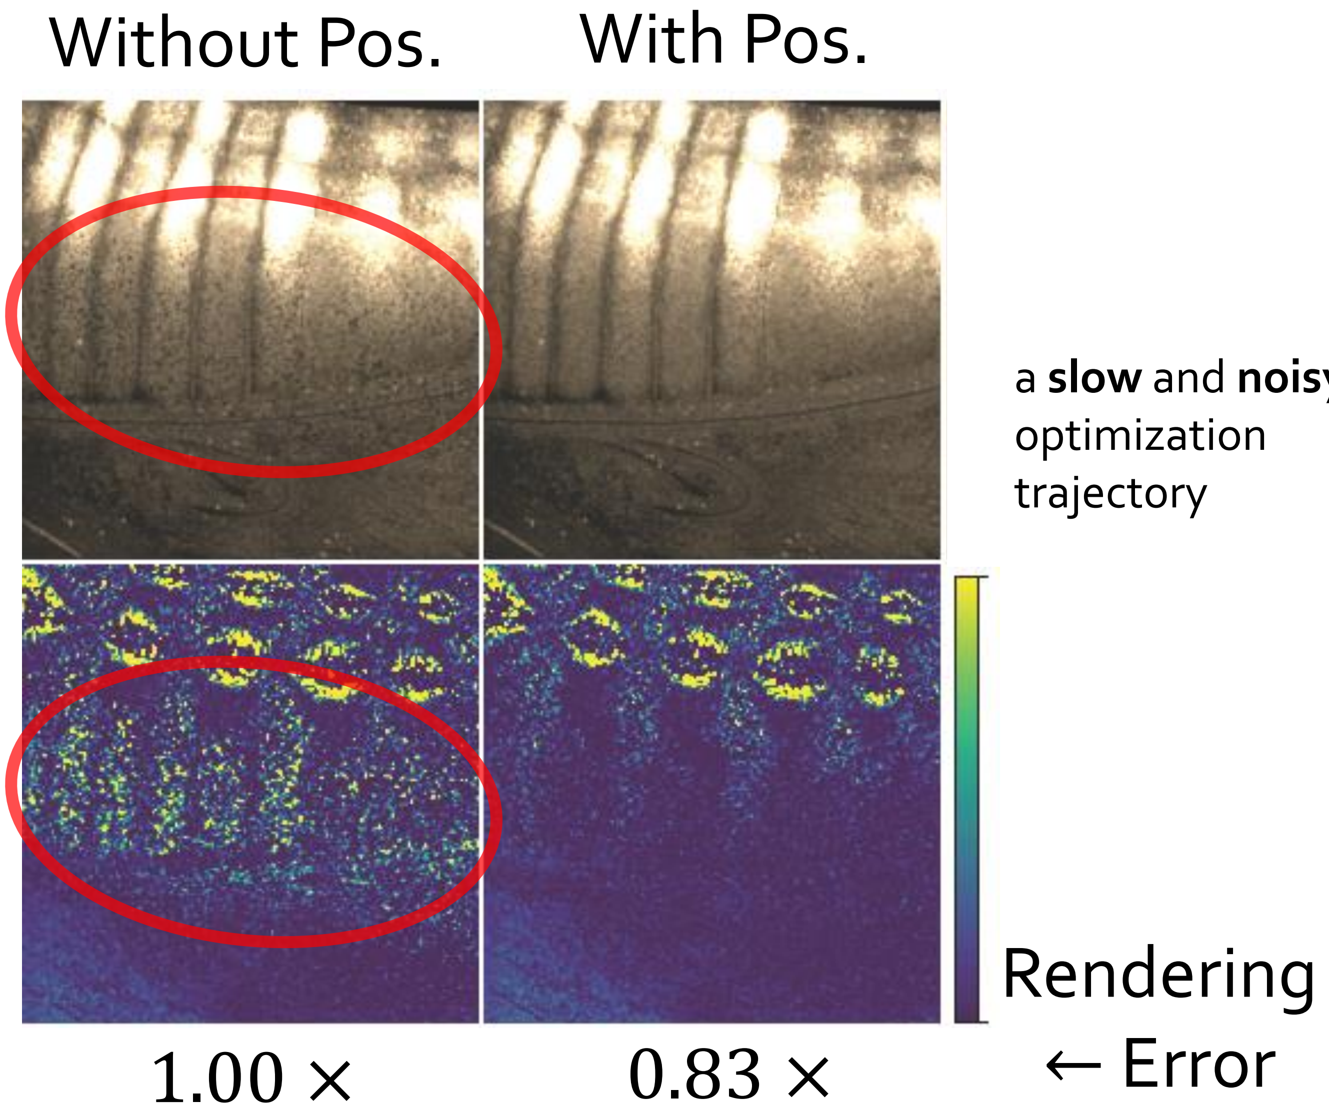

a **slow** and **noisy** optimization trajectory

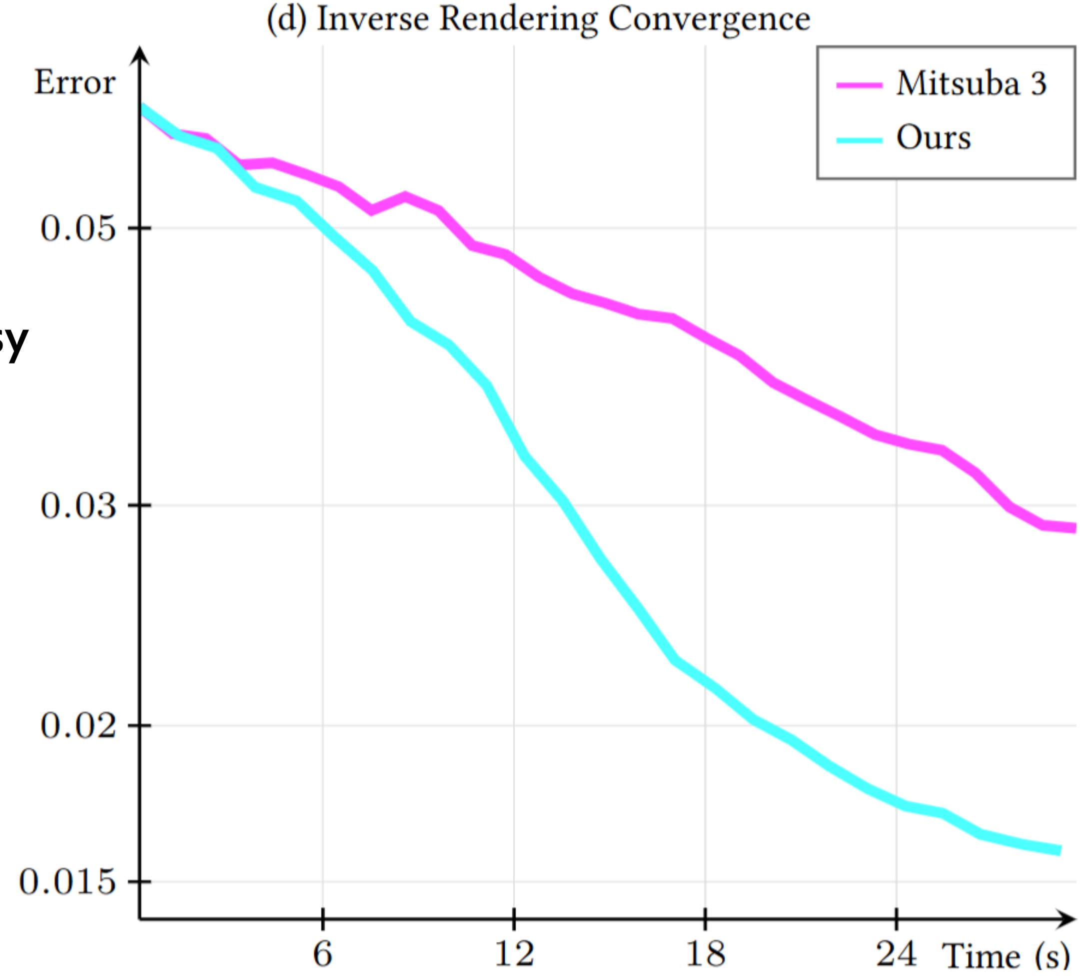

Results: Inverse Rendering Video – 1 spp

Initial

Target

Mitsuba 3 Baseline

Ours

Summary

» **Parameter-space differentiable rendering** enables efficient derivative reuse.

» **Positivized RIS** extends RIS to real-valued functions to achieve theoretical zero-variance convergence of resampled derivative estimates.

» **Reusing samples** from previous gradient descent iterations results in faster inverse rendering.

» **Limitation:**

- ➢ Assumes gradient correlation, which may fail at high learning rates.
- ➢ Impact of gradient errors on convergence speed is unclear.

Conclusion

» Physically-based differentiable rendering has historically been slow.

» But we can leverage decades of (real-time) rendering research to make it fast.

» Our framework is applicable to other optimization problems outside rendering.

Quiz