Team 1

Mid-term Project Presentation

Adding indirect lighting to ReSTIR DR

Niklas Sanden
Tan Chao

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

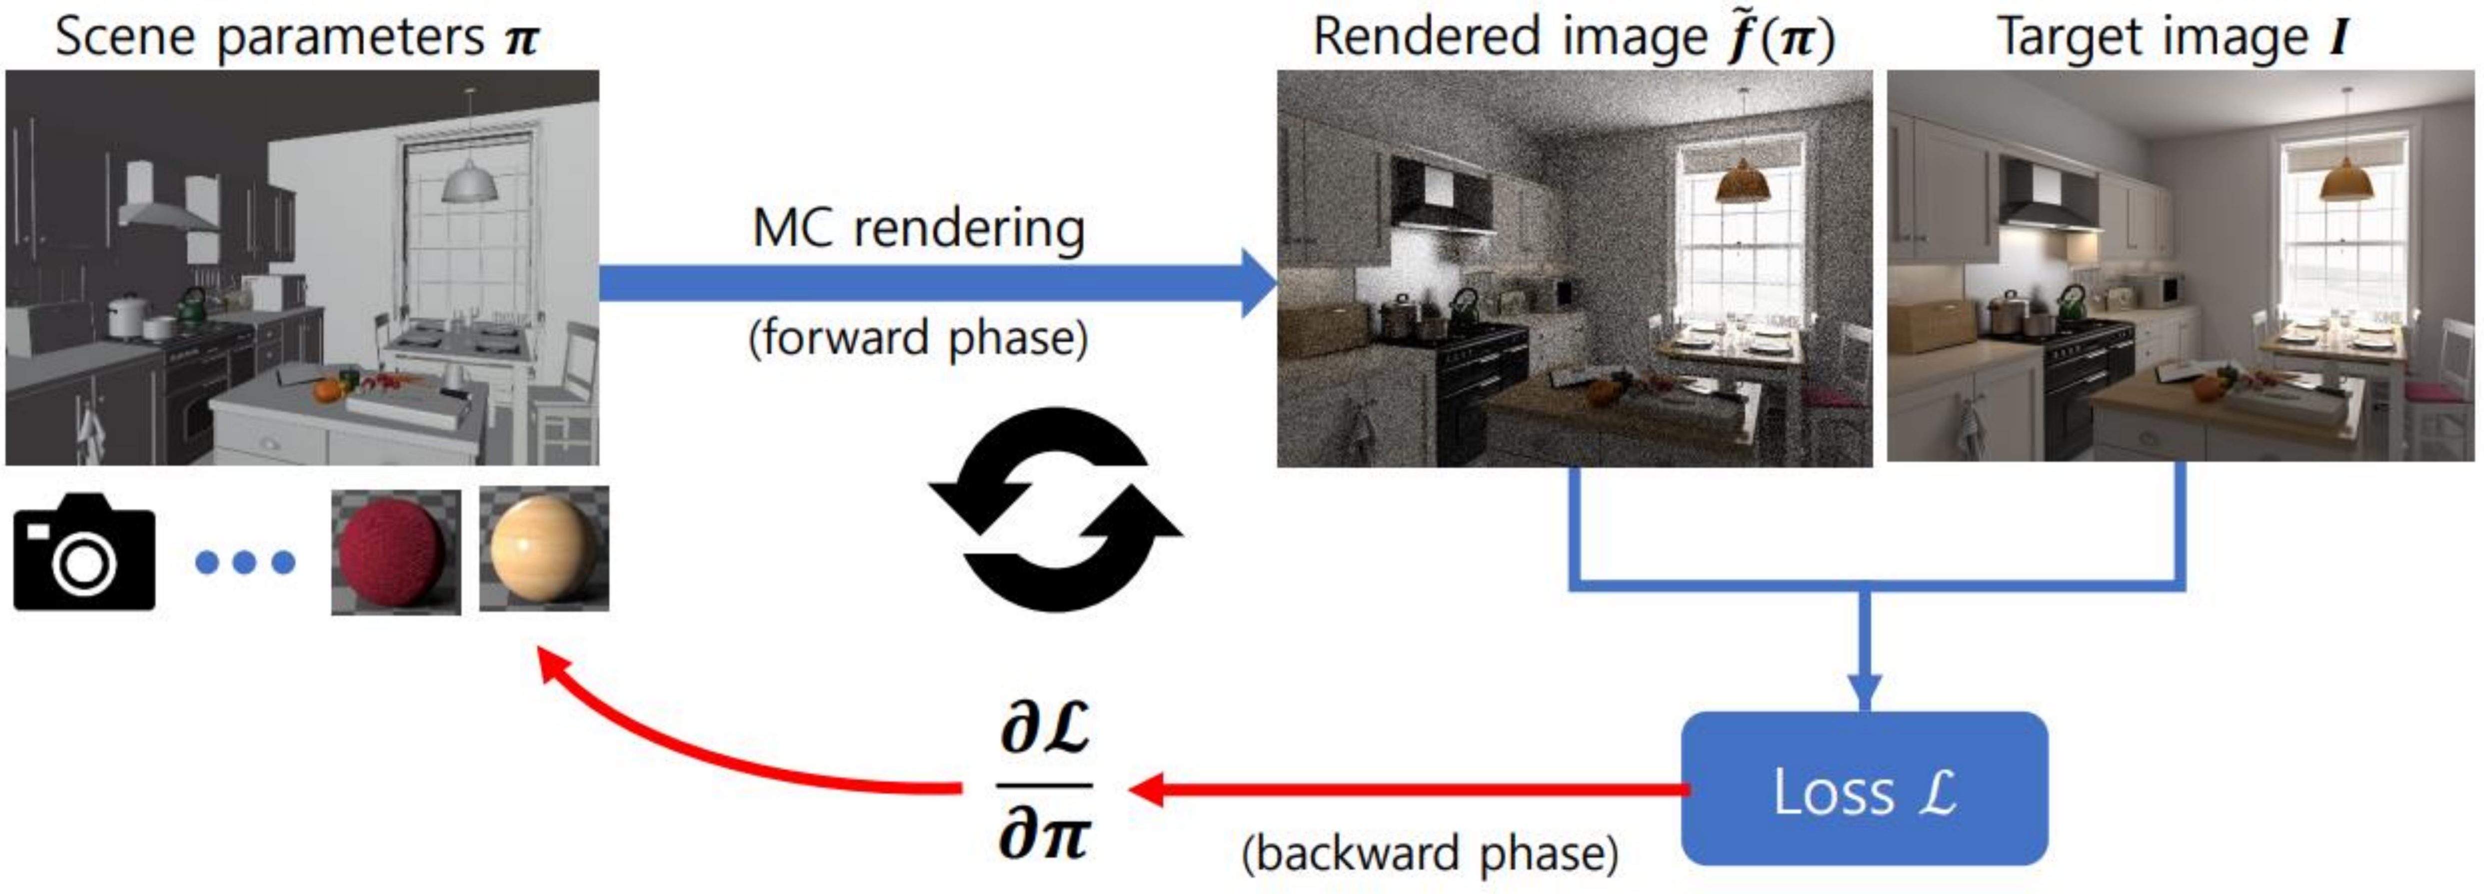

Target-Aware Image Denoising for Inverse Monte Carlo Rendering, Bochang Moon

Inverse rendering (materials)

Initial

Target

Noisy gradients

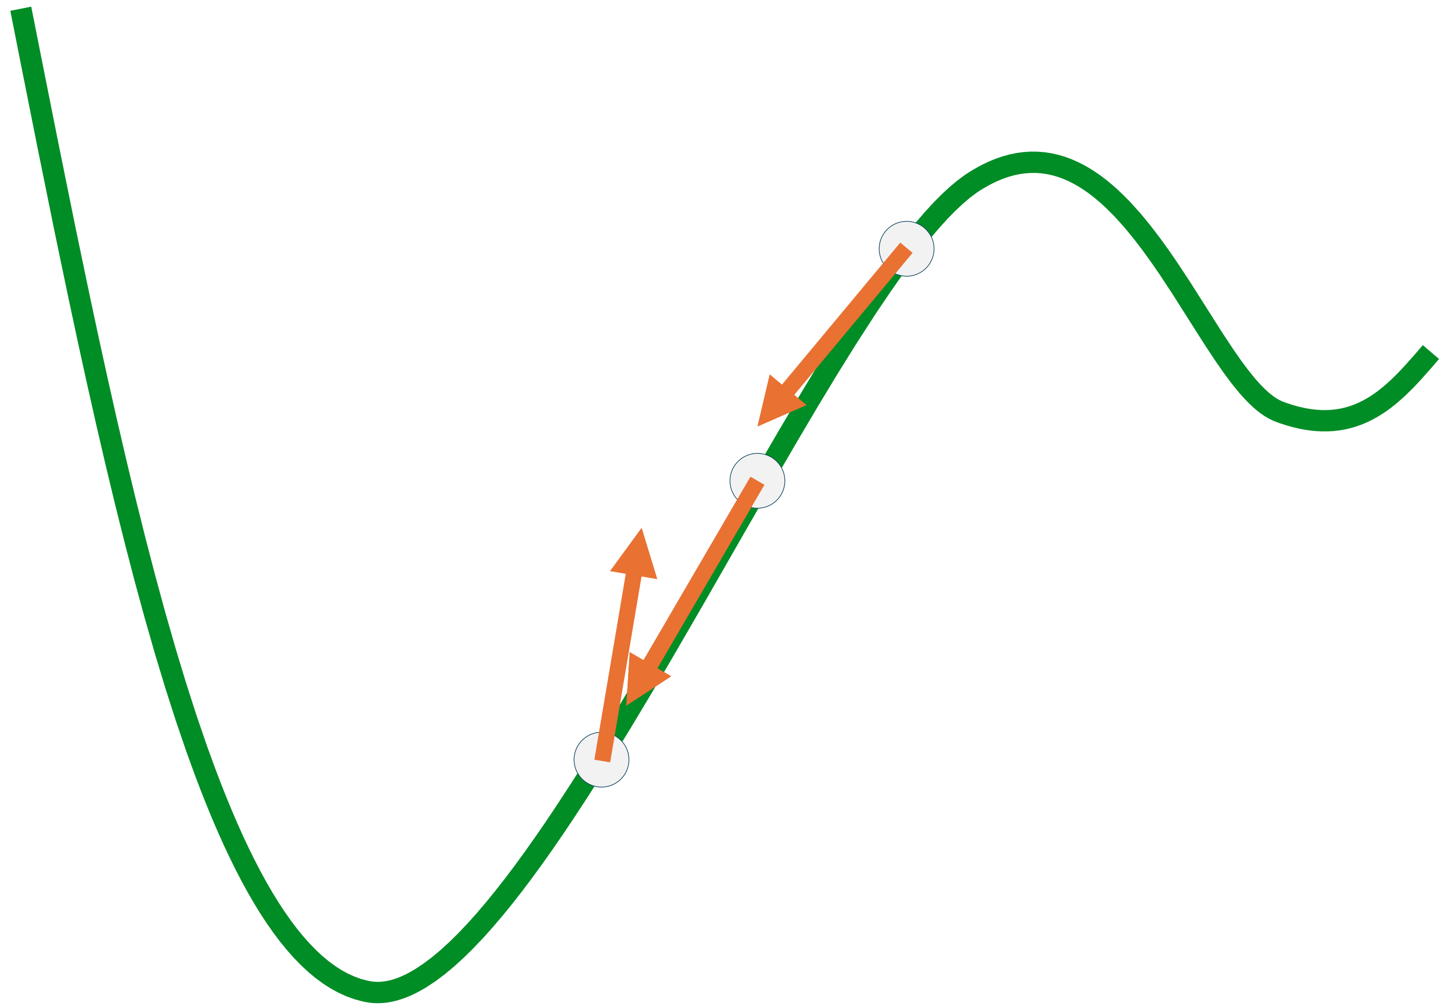

Bitterli et al. *SIGGRAPH 2020*. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting.

ReSTIR (unbiased spatiotemporal reuse)

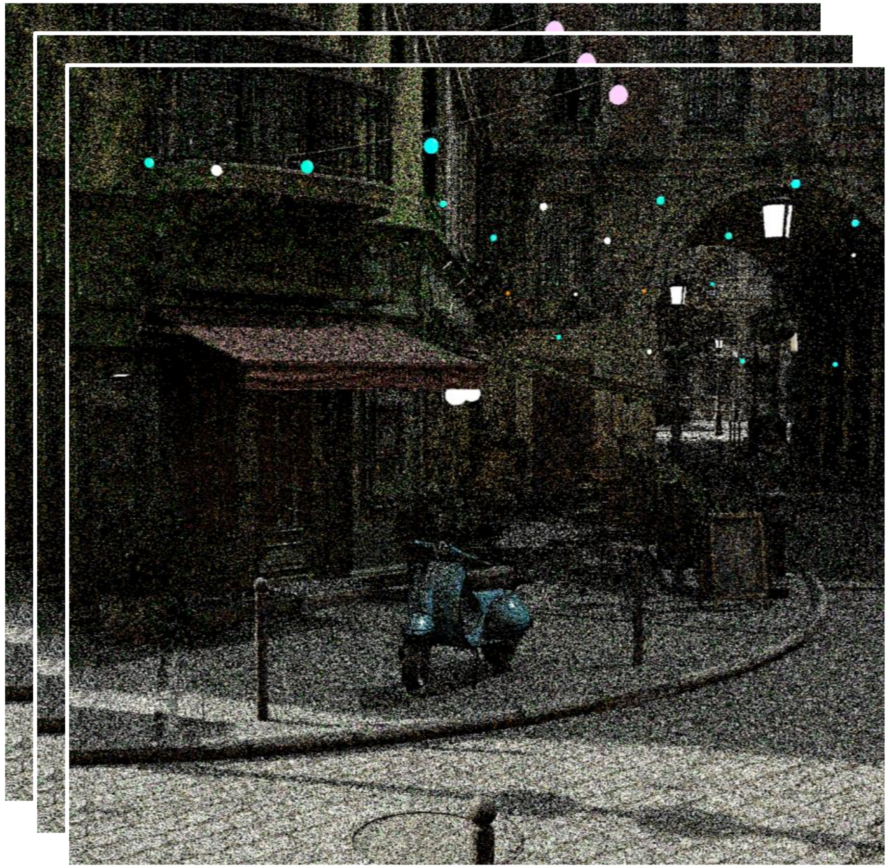

Sequence of similar noisy frames

Reuse of previous frames

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)

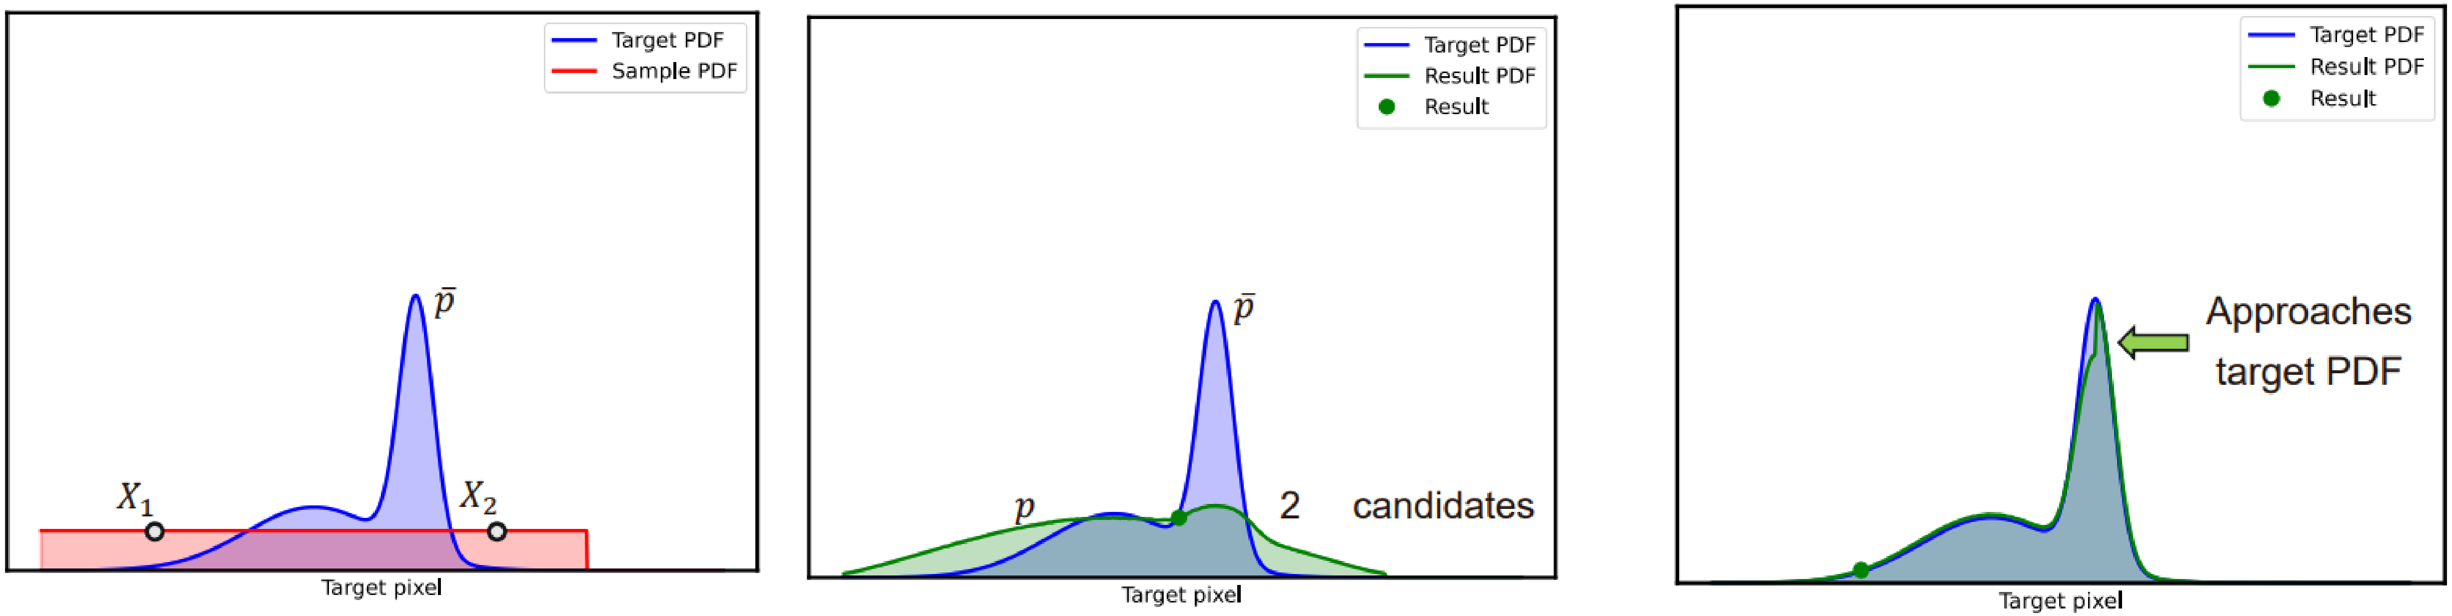

A Gentle Introduction to ReSTIR: Path Reuse in Real-time, Wyman et al. *ACM SIGGRAPH 2023 Courses*

ReSTIR temporal reuse

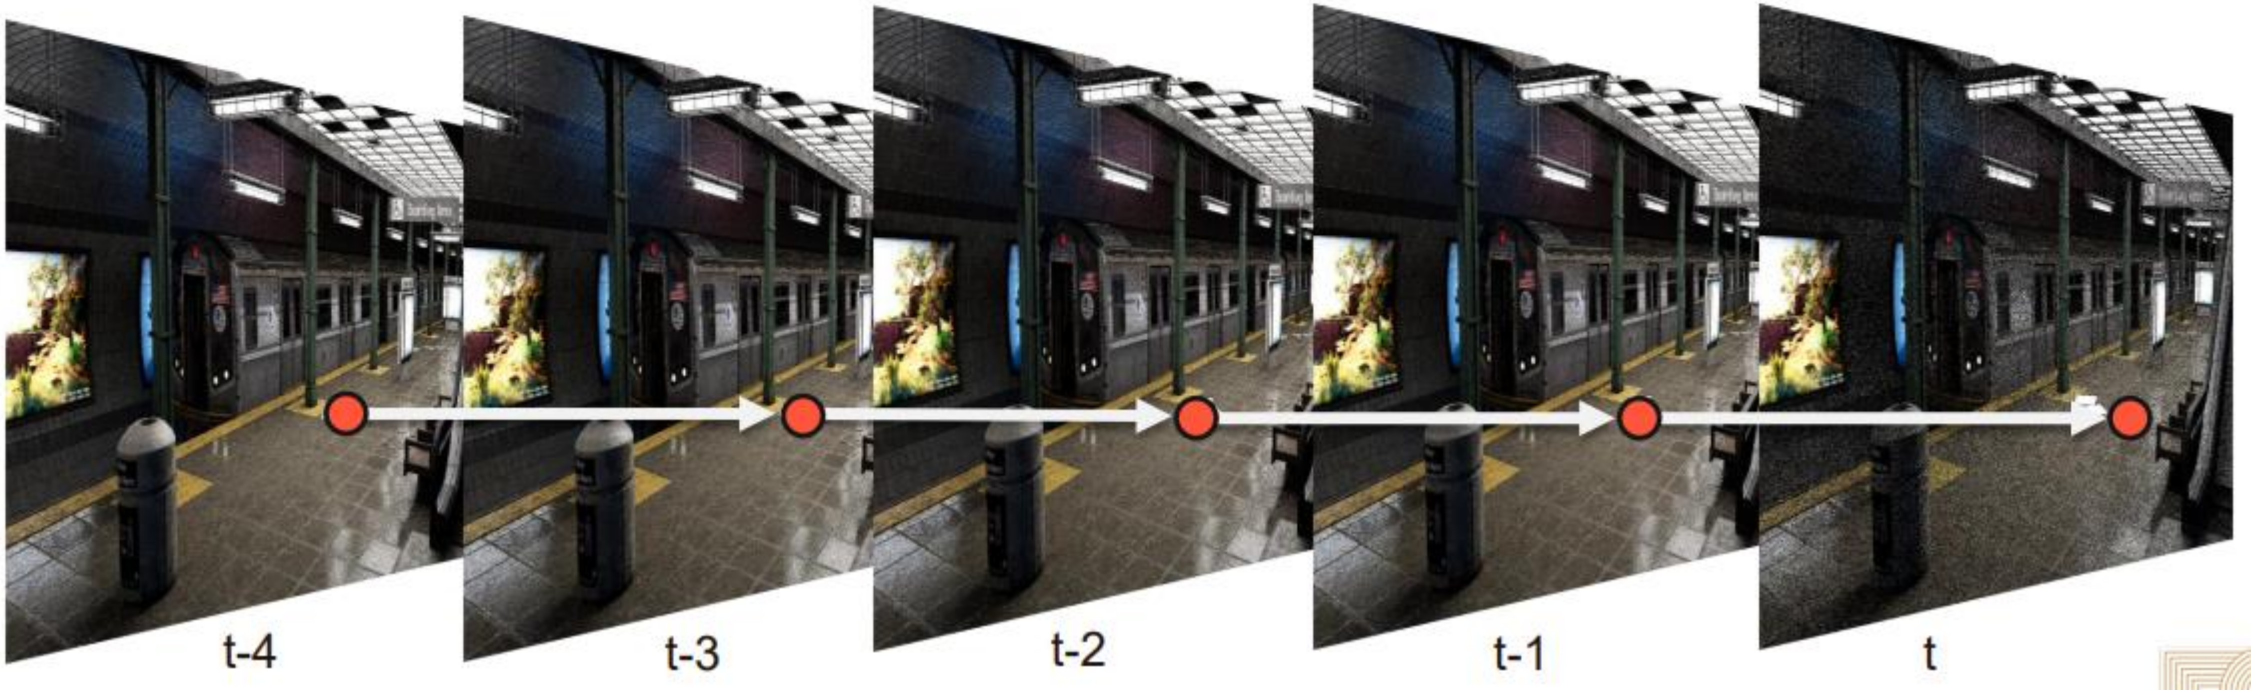

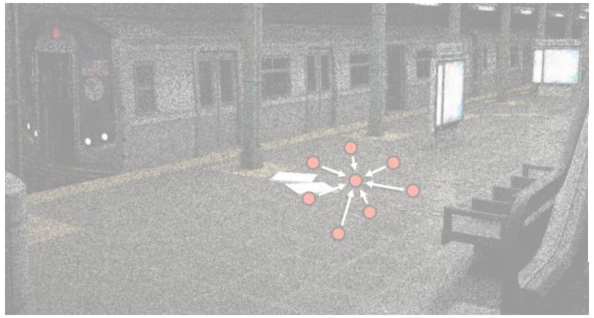

(spatial reuse is not used in ReSTIR DR)

A Gentle Introduction to ReSTIR: Path Reuse in Real-time, Wyman et al. *ACM SIGGRAPH 2023 Courses*

Theoretical contributions of the paper

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

- This will not be modified in our project

The Problem with Pixel-centric Differentiable Rendering

Forward Rendering

Single intensity I for each of N pixels = N samples

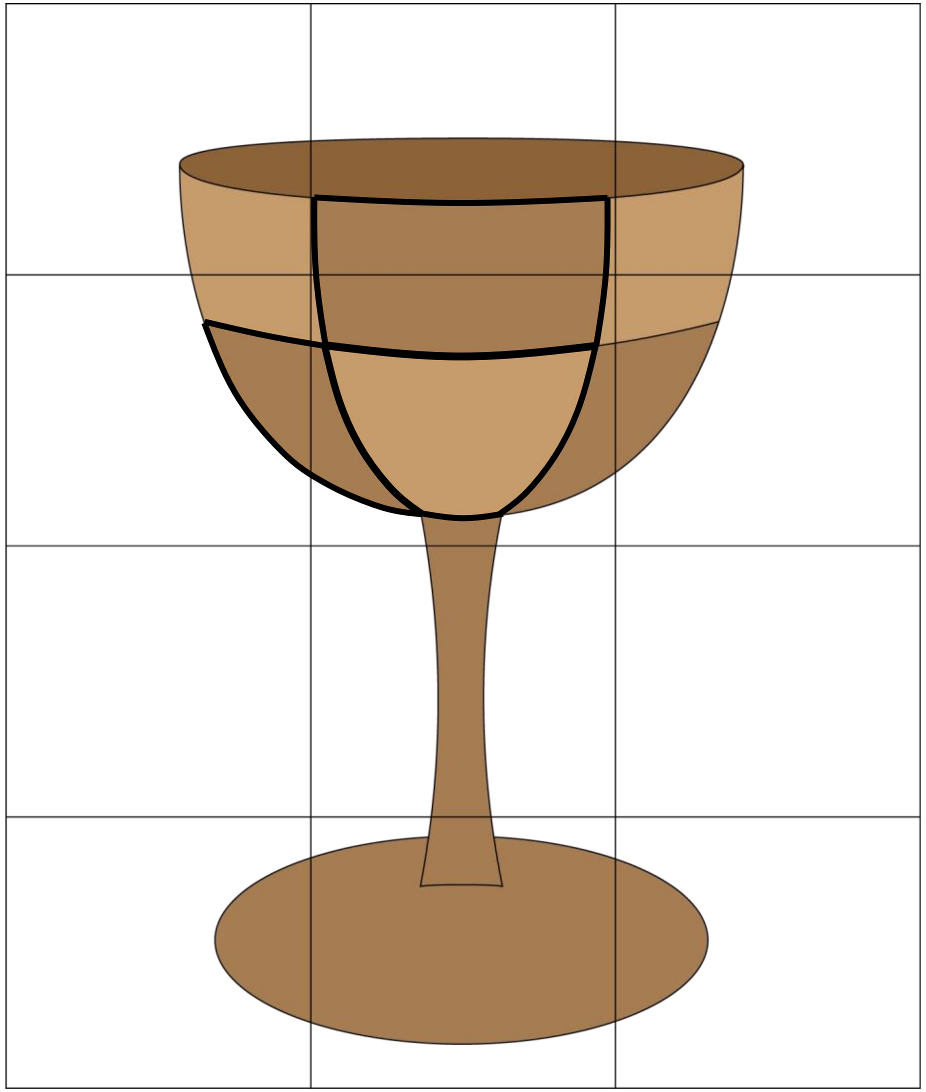

Differentiable Rendering

One derivative for each texel π_i in each pixel

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

M texels = N · M samples

Texture Optimization Algorithm

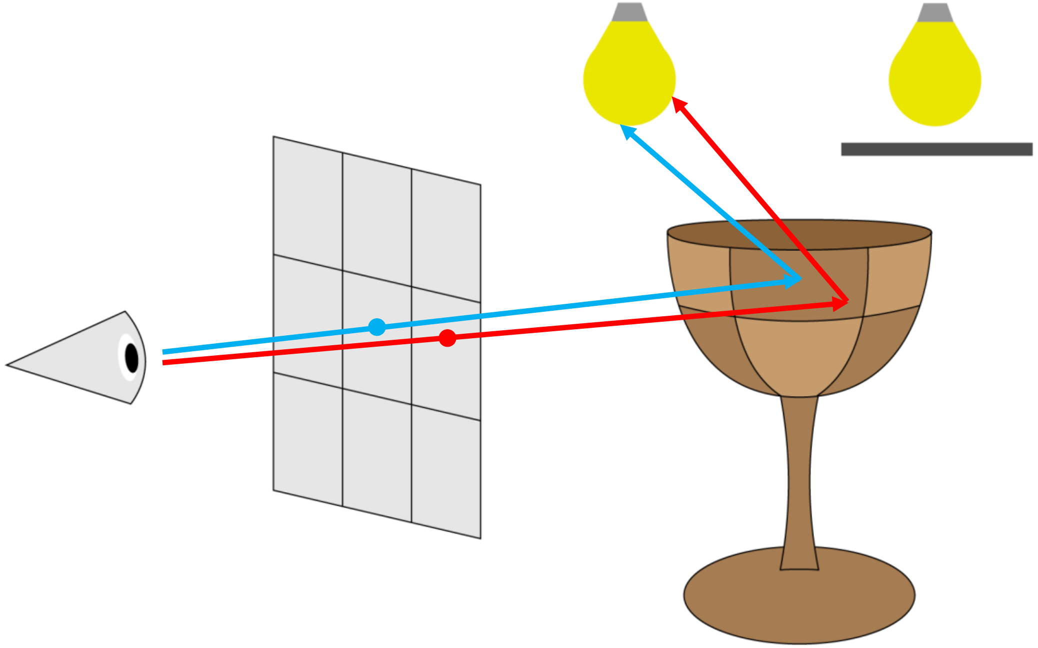

Store
1 positive, 1 negative
sample per texel

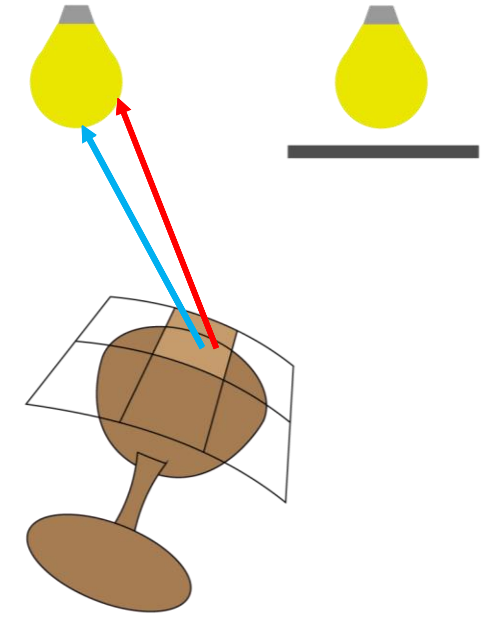

Iteration i

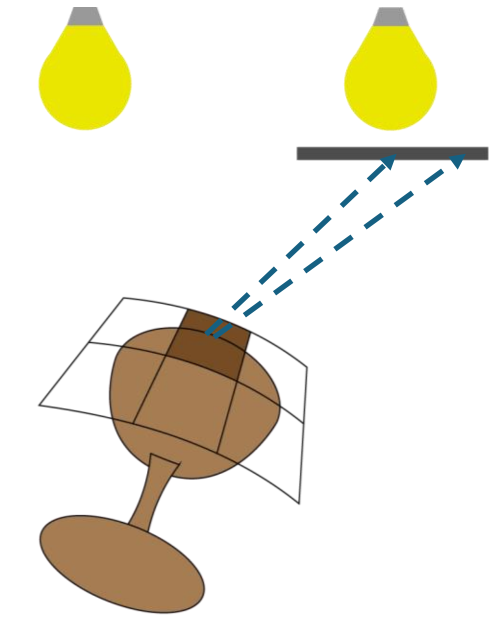

Iteration $i + 1$

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.

Project

Future work. Reuse across parameters, analogous to spatial reuse in ReSTIR, is possible with our parameter-space formulation. A potential challenge lies in efficiently selecting neighboring parameters to reuse. Reuse introduces correlation in sample estimates [Sawhney et al. 2022], and the exact effect of correlated gradients in inverse optimization is an interesting avenue to be investigated. While we have focused on differentiable and inverse rendering for BRDF textures under complex direct lighting, our theory and methods are general and can be extended to other rendering methods and scenarios involving general light transport, discontinuities, and other parameters, such as in volumetric or neural representations. Finally, our PGRIS estimator is immediately applicable to Monte Carlo integral estimation in contexts broader than rendering, where integrands can be both positive and negative.

Motivation

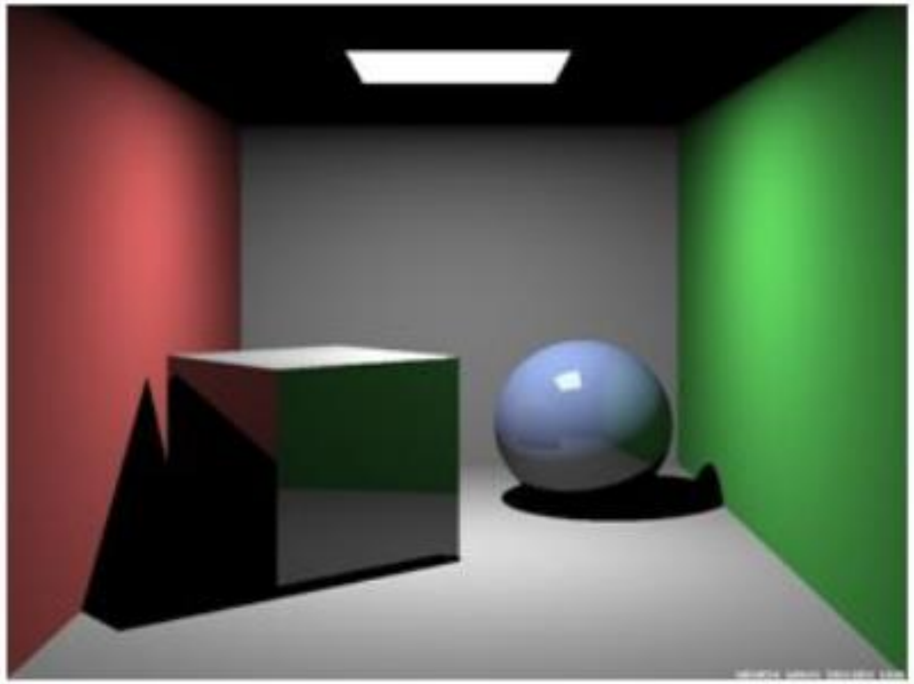

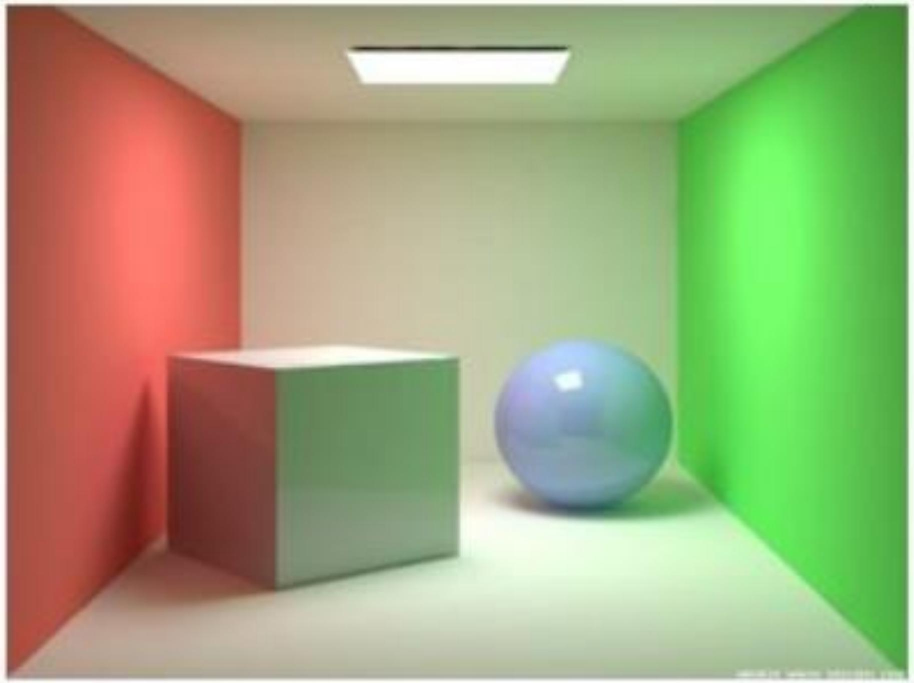

Ouyang et al. HPG *2021*.
ReSTIR GI: Path Resampling for Real-Time Path Tracing

ReSTIR GI

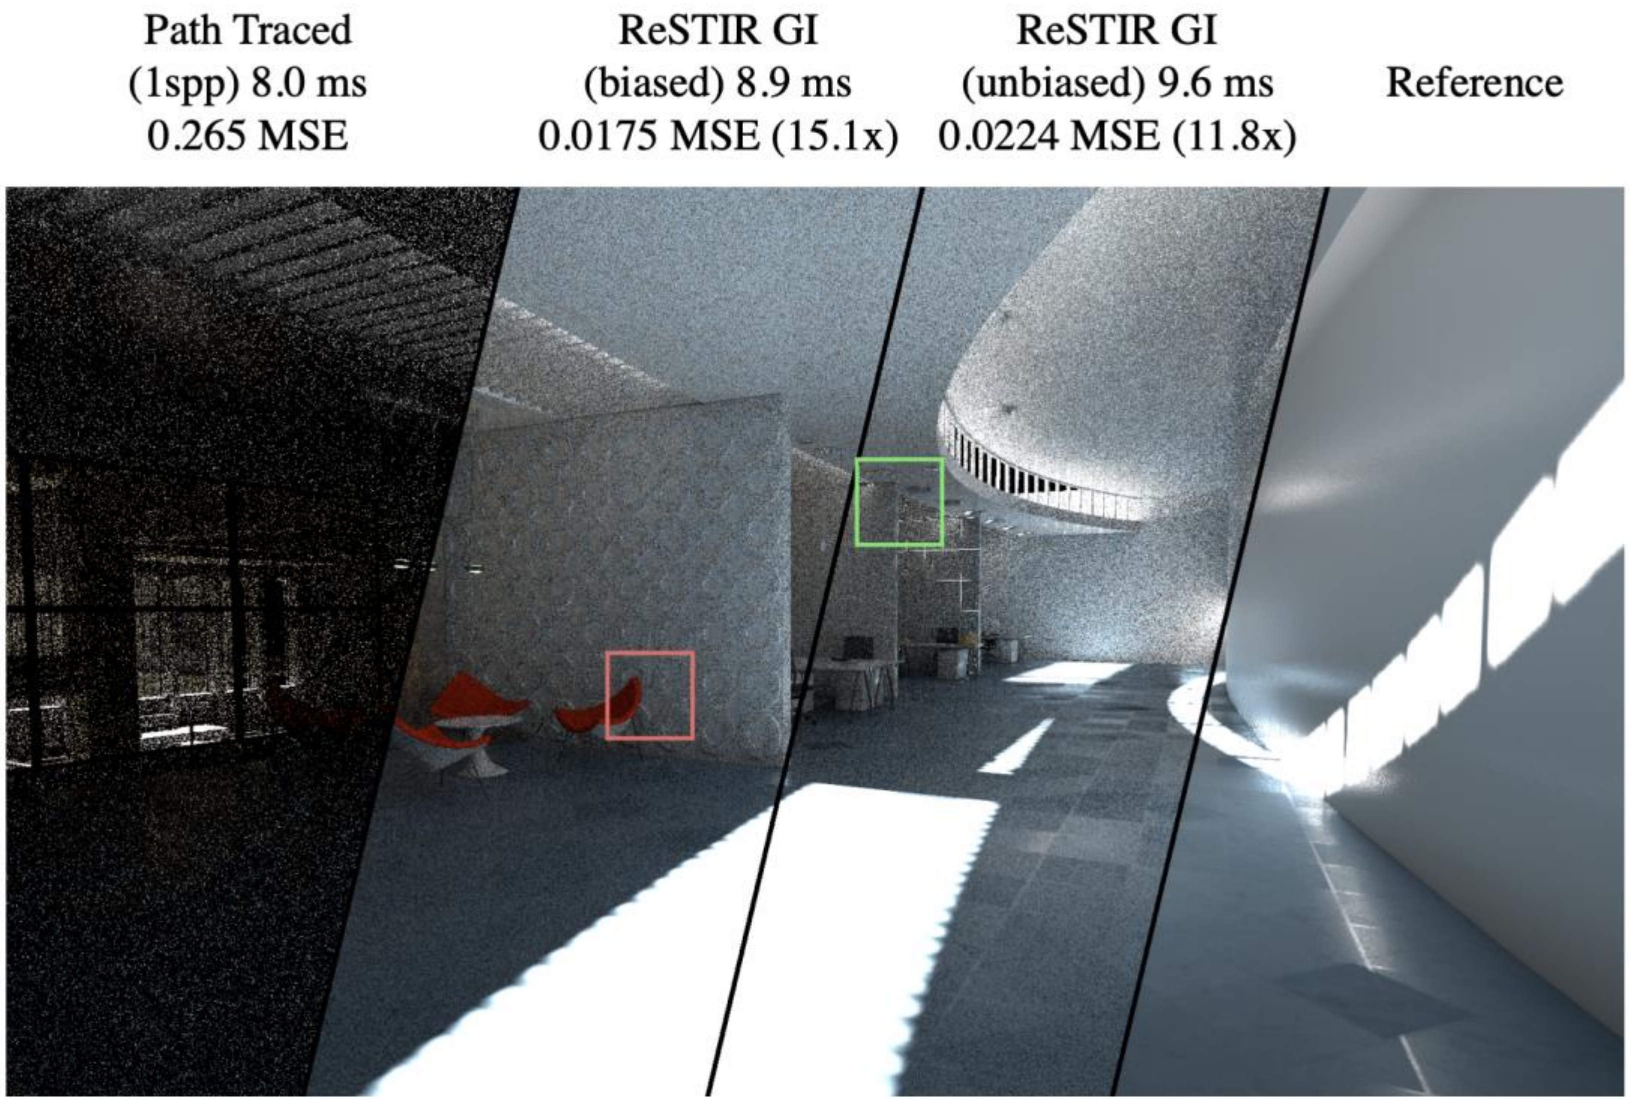

Multiple bounces

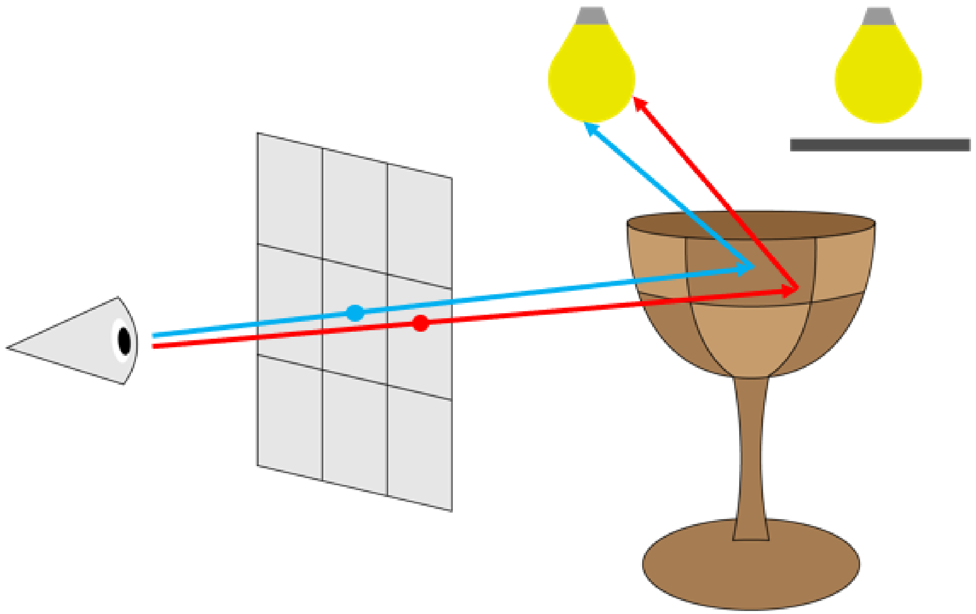

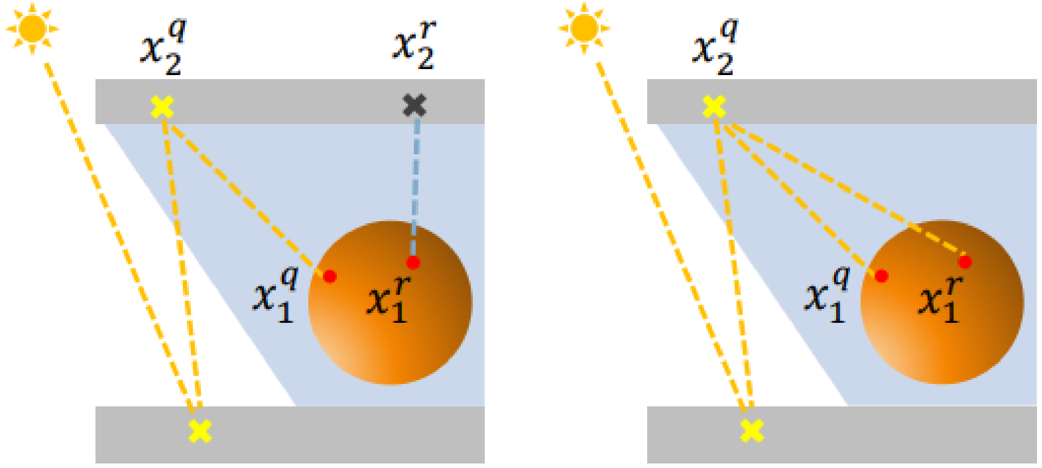

TODO

- Mitsuba 3
- Implement ReSTIR GI
- Test scenes/experiments

Roles

- Niklas Sanden
  - Try to compile Mitsuba
  - Extend ReSTIR DR with GI

- Tan Chao
  - Try compile Mitsuba
  - Try to find different scenes for indirect illumination
    - Set the experiments up in Mitsuba