SIGGRAPH 2024
DENVER+ 28 JUL — 1 AUG

THE PREMIER CONFERENCE & EXHIBITION ON COMPUTER GRAPHICS & INTERACTIVE TECHNIQUES

Conditional Mixture Path Guiding for Differentiable Rendering

Zhimin Fan Pengcheng Shi Mufan Guo
Ruoyu Fu Yanwen Guo Jie Guo

Nanjing University, China

Differentiable Rendering

- Compute **derivatives** of pixel intensity w.r.t. scene parameters

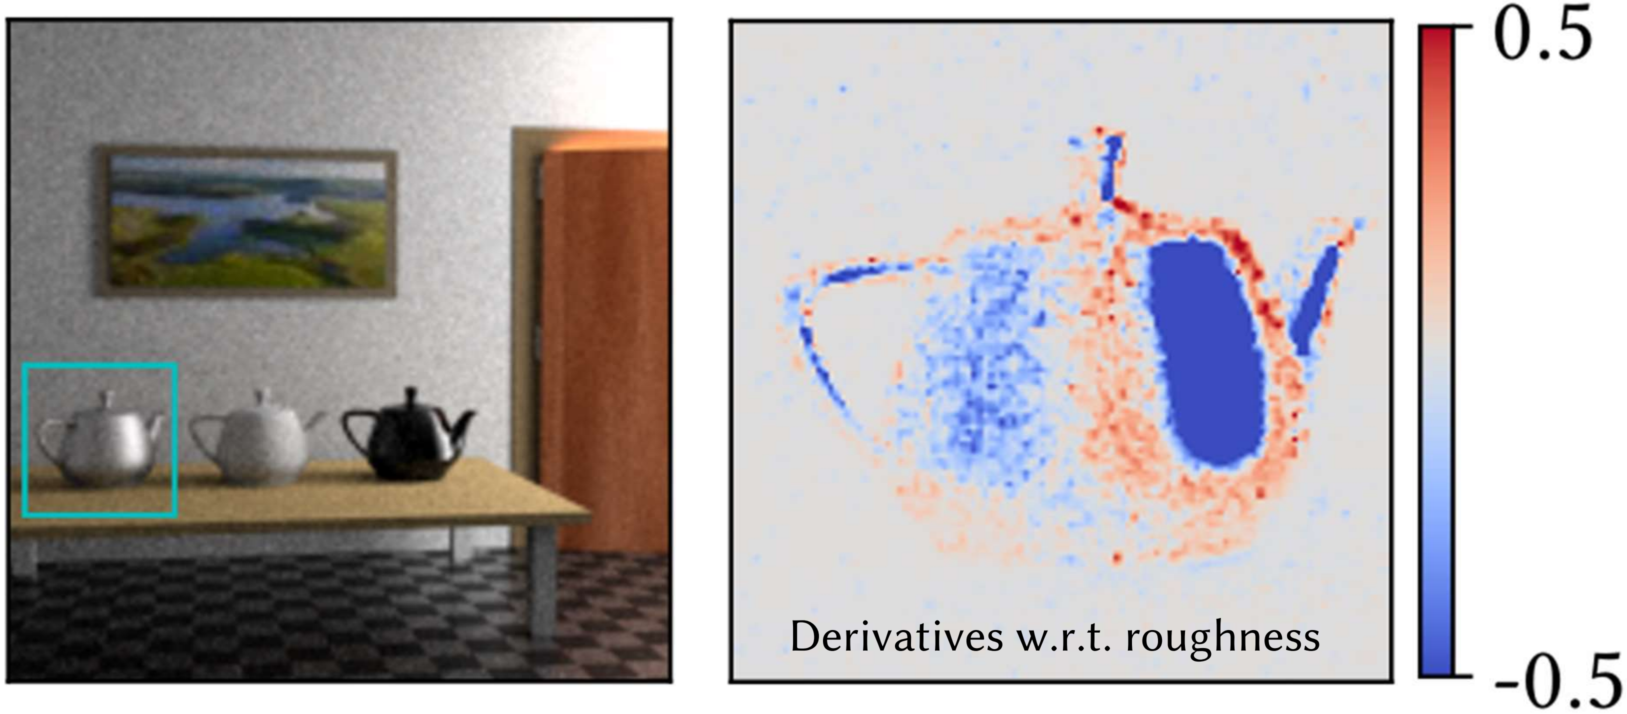

Differentiable Rendering

- Compute **derivatives** of pixel intensity w.r.t. scene parameters
- Enable gradient-based optimization for **inverse** reconstruction

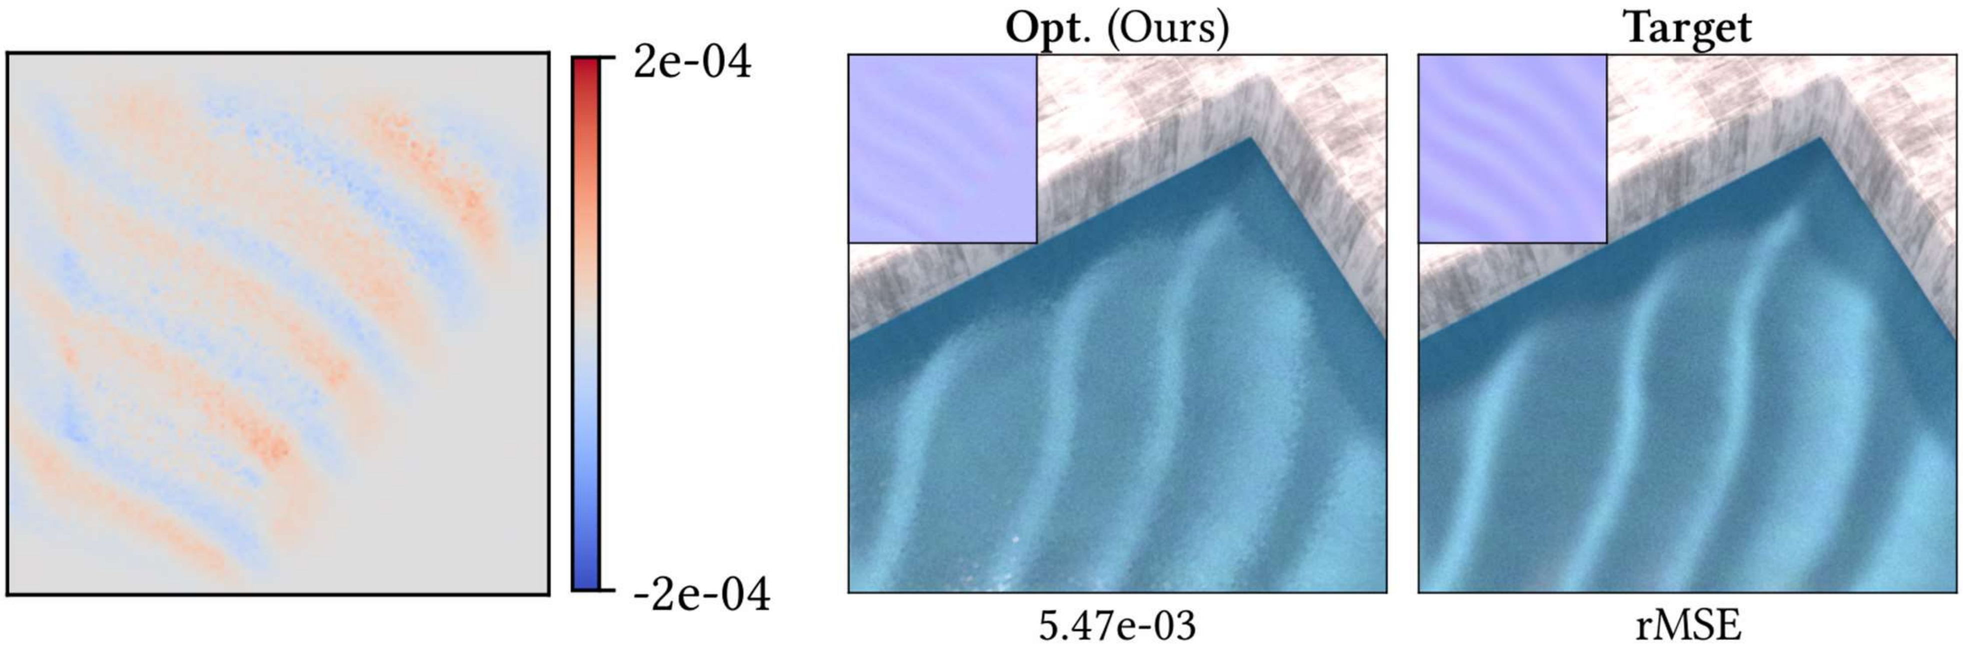

Differential Rendering Equation

$$\partial_\pi L_o(\boldsymbol{x}, \boldsymbol{\omega}_o) = \partial_\pi L_e(\boldsymbol{x}, \boldsymbol{\omega}_o) + \int_{\mathcal{S}^2} \left(f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \partial_\pi L_i(\boldsymbol{x}, \boldsymbol{\omega}_i) + \partial_\pi f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\boldsymbol{x}, \boldsymbol{\omega}_i) \right) \mathrm{d}\boldsymbol{\omega}_i^\perp$$

- Differential radiance that is emitted from light sources
- Differential radiance that scatters like ordinary radiance
- Differential radiance that is also added on the shading point with differentiable BSDF

Importance sampling

- Monte Carlo estimators using samples from $p(x)$

$$\langle I \rangle = \left\langle \int_{\Omega} f(x)\,\mathrm{d}x \right\rangle = \frac{1}{N} \sum_{j=1}^{N} \frac{f(X_j)}{p(X_j)}$$

- The shape of $p(x)$ should follow the shape of $f(x)$
- Variance reaches zero if $p(x) \propto f(x) \quad (f(x) > 0)$
- We consider single-signed functions first

Important paths for material derivatives

A path with non-zero contribution must connect

- objects with differentiable params
- light sources

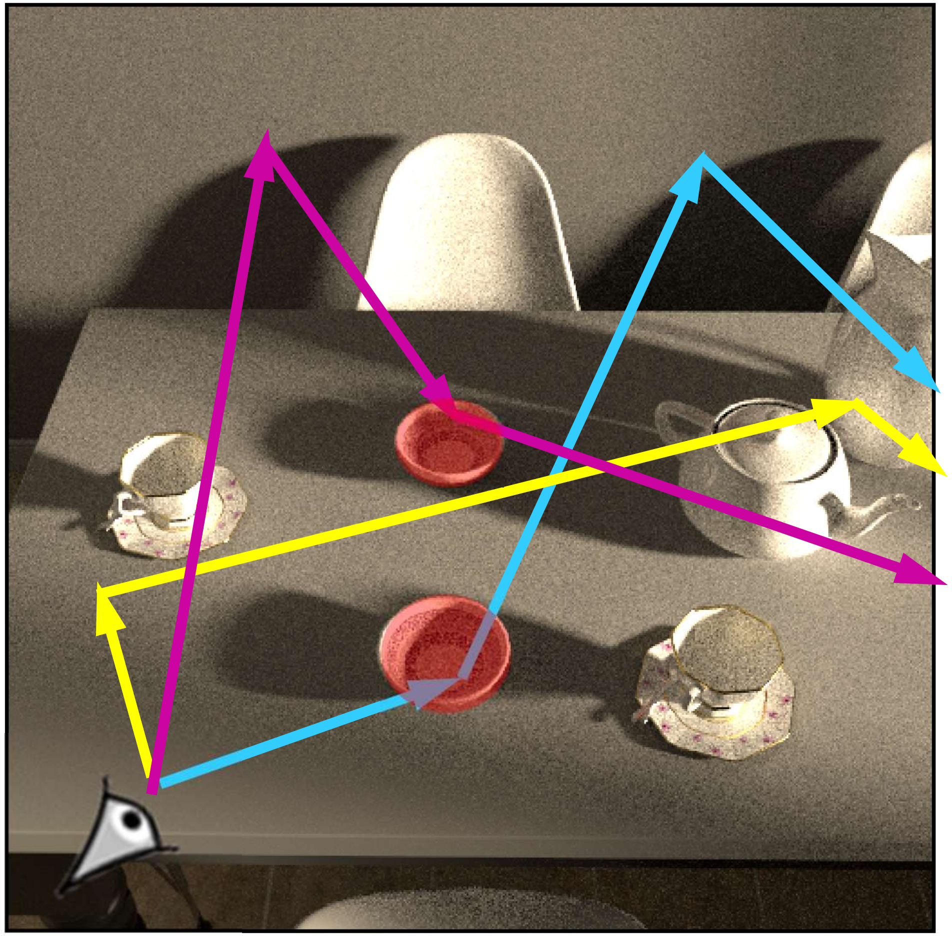

In this figure:

- The yellow path has zero contribution
- The blue and magenta paths are good

Sampling techniques for forward rendering

- Sampling w.r.t. product of incident radiance and BSDF

$$p^L(\omega_i|\boldsymbol{x}, \omega_o) \propto f_s(\boldsymbol{x}, \omega_i, \omega_o) L_i(\boldsymbol{x}, \omega_i)$$

- Clear failure cases

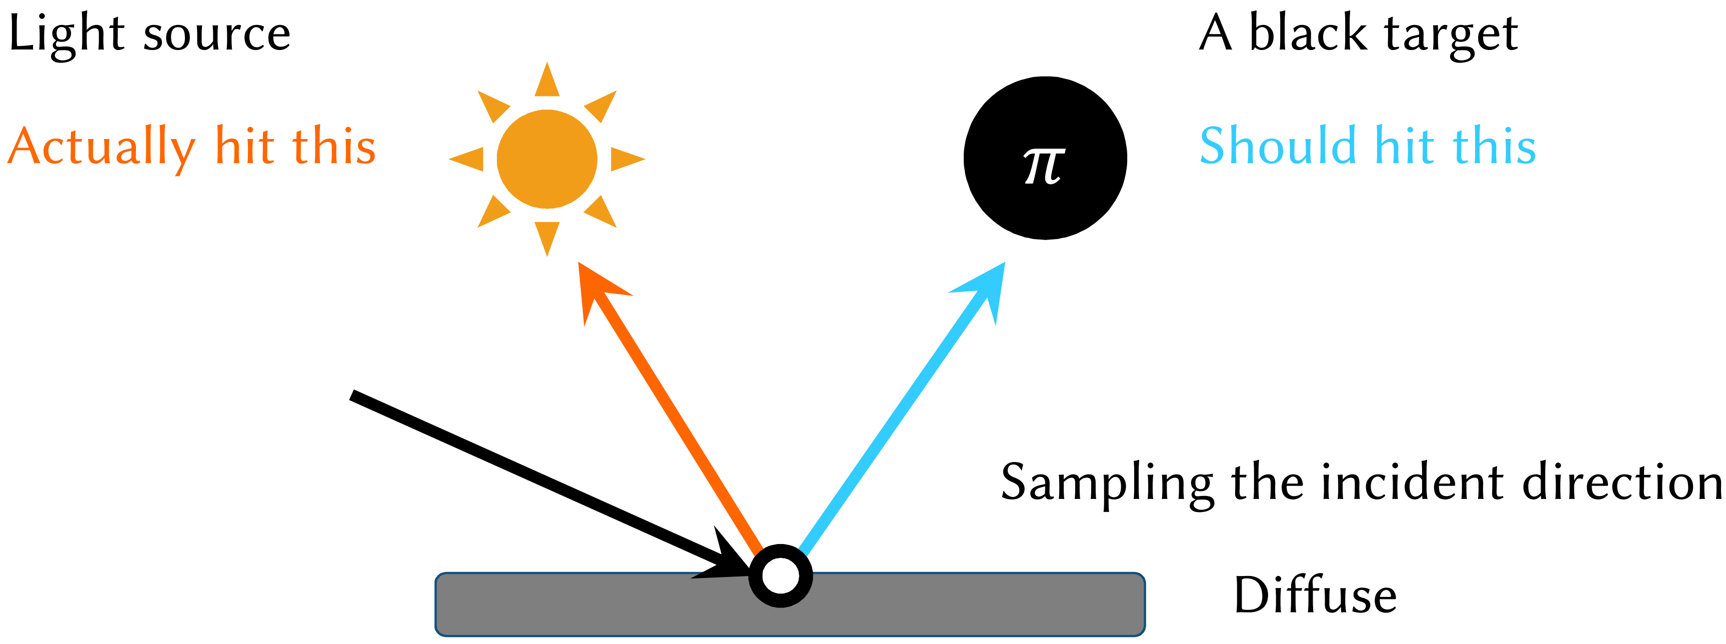

Local differential sampling

- Consider the DRE

$$\partial_\pi L_o(\boldsymbol{x}, \boldsymbol{\omega}_o) = \partial_\pi L_e(\boldsymbol{x}, \boldsymbol{\omega}_o) + \int_{\mathcal{S}^2} \left(f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \partial_\pi L_i(\boldsymbol{x}, \boldsymbol{\omega}_i) + \partial_\pi f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\boldsymbol{x}, \boldsymbol{\omega}_i) \right) \mathrm{d}\boldsymbol{\omega}_i^\perp$$

- Sample proportional to the integrand

$$p^D(\boldsymbol{\omega}_i | \boldsymbol{x}, \boldsymbol{\omega}_o) \propto f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \partial_\pi L_i(\boldsymbol{x}, \boldsymbol{\omega}_i) + \partial_\pi f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\boldsymbol{x}, \boldsymbol{\omega}_i)$$

- Good for direct illumination, but ...

Challenge: paths are shared

- $\partial_\pi L_i$ requires a recursive estimation of the L_i in the term $\partial_\pi f_s(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\boldsymbol{x}, \boldsymbol{\omega}_i))$
- Modern practices share the same set of path samples for $\partial_\pi L_i$ and L_i
- p^D does not always achieve the goal of importance sampling

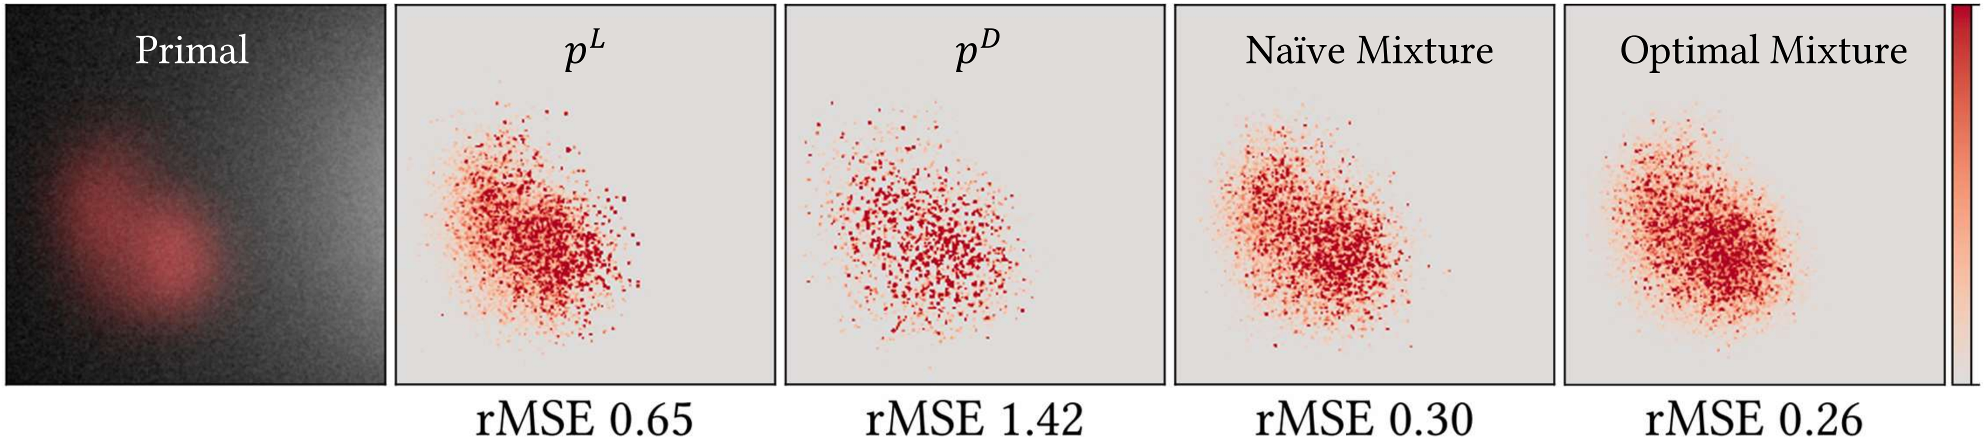

- Using only p^L or p^D to sample paths could lead to high variance

Mixture sampling

- A straightforward solution is to use a **mixture** of p^L and p^D

$$p^M(\boldsymbol{\omega}_i|\boldsymbol{x},\boldsymbol{\omega}_o) \propto w^L f_s(\boldsymbol{x},\boldsymbol{\omega}_i,\boldsymbol{\omega}_o)L_i(\boldsymbol{x},\boldsymbol{\omega}_i)+$$
$$w^D\left(f_s(\boldsymbol{x},\boldsymbol{\omega}_i,\boldsymbol{\omega}_o)\partial_\pi L_i(\boldsymbol{x},\boldsymbol{\omega}_i)+\partial_\pi f_s(\boldsymbol{x},\boldsymbol{\omega}_i,\boldsymbol{\omega}_o)L_i(\boldsymbol{x},\boldsymbol{\omega}_i)\right)$$

- How to determine w^L and w^D?

Optimal mixture

- Given a path prefix $\boldsymbol{x}_1, \boldsymbol{x}_2, \dots, \boldsymbol{x}_i$
- When sampling the direction from $\boldsymbol{x}_i$ to $\boldsymbol{x}_{i+1}$
- The contribution of the entire path writes

$$\partial_\pi \left(L_o \prod_{j=1}^{i-1} f_j \right) = L_o \, \partial_\pi \prod_{j=1}^{i-1} f_j + (\partial_\pi L_o) \prod_{j=1}^{i-1} f_j$$

- Thus the optimal weight

$$w_i^L(\overleftarrow{\boldsymbol{x}_i}) = \partial_\pi \prod_{j=1}^{i-1} f_j, \quad w_i^D(\overleftarrow{\boldsymbol{x}_i}) = \prod_{j=1}^{i-1} f_j$$

Reduction of dimensionality

$$p(\omega_i|\overleftarrow{\boldsymbol{x}_i}) \propto w^L(\overleftarrow{\boldsymbol{x}_i}) f_s(\boldsymbol{x}, \omega_i, \omega_o) L_i(\boldsymbol{x}, \omega_i) +$$
$$w^D(\overleftarrow{\boldsymbol{x}_i}) \left(f_s(\boldsymbol{x}, \omega_i, \omega_o) \partial_\pi L_i(\boldsymbol{x}, \omega_i) + \partial_\pi f_s(\boldsymbol{x}, \omega_i, \omega_o) L_i(\boldsymbol{x}, \omega_i) \right)$$

$$w_i^L(\overleftarrow{\boldsymbol{x}_i}) = \partial_\pi \prod_{j=1}^{i-1} f_j, \quad w_i^D(\overleftarrow{\boldsymbol{x}_i}) = \prod_{j=1}^{i-1} f_j$$

- p is conditioned on the path prefix, which could be high-dimensional
- We instead fit p^D and p^L separately and mix them on the fly

Summary: conditional mixture sampling

A **Unidirectional** Method for Importance Sampling Path **Derivatives**

Primal and Differential **Distributions**

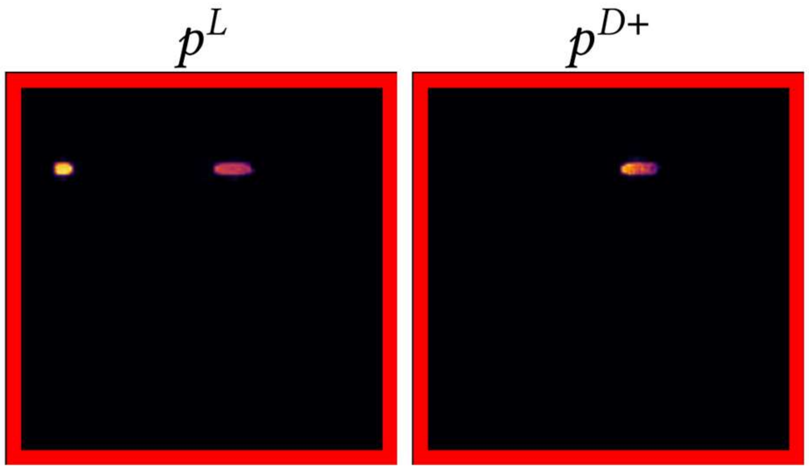

$$p \propto w^L \hat{p}^L + w^D \hat{p}^D$$

$\hat{p}$ refers to an unnormalized distribution

Analytic Mixture **Weights** Conditioned on Path Prefixes

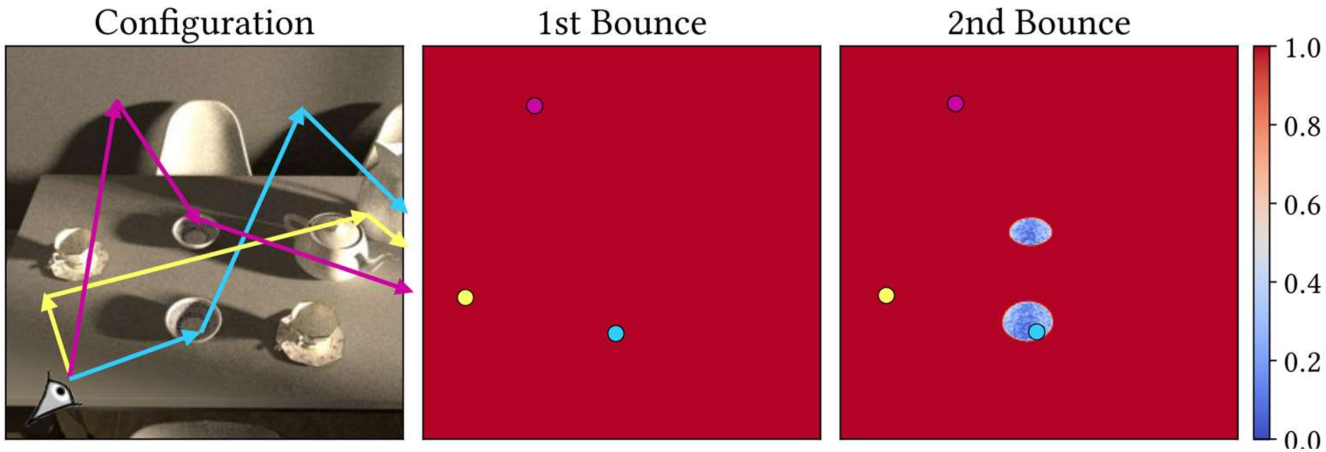

$$w_i^L(\overleftarrow{x_i}) = \partial_\pi \prod_{j=1}^{i-1} f_j, \quad w_i^D(\overleftarrow{x_i}) = \prod_{j=1}^{i-1} f_j$$

Application with path guiding

Path guiding

- Fit distributions from historical samples
- Target at L_i or $f_s L_i$

Conditional mixture path guiding

- Fit distributions for $p^L \propto f_s L_i$ and $p^D \propto \partial_\pi(f_s L_i)$
- Samples are from the previous optimization steps
- Estimate the mean μ^L and μ^D for two target distributions, respectively
- Compute the mixture on-the-fly

$$p \propto w^L \mu^L p^L + w^D \mu^D p^D$$

More details

Please refer to the paper for

- Positivization for sign-variance elimination
- Extension to multiple parameters (L1 norm of gradients)
- Distribution model using kd-trees and quadtrees
- Distribution sharing across optimization steps

$$\frac{f^+(X^+)}{p^+(X^+)} + \frac{f^-(X^-)}{p^-(X^-)}$$

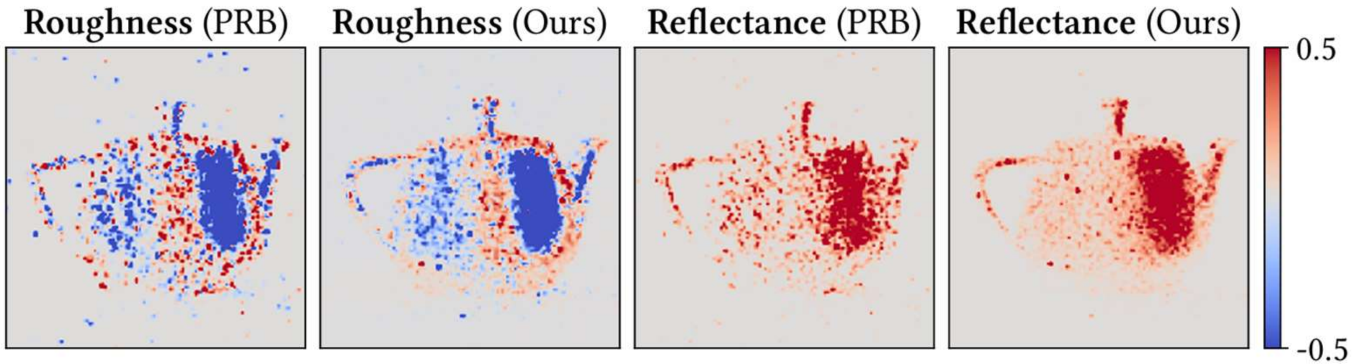

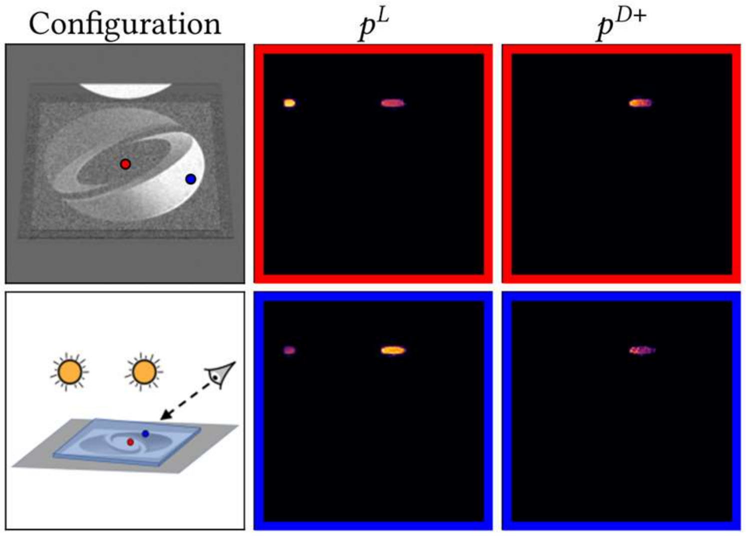

Results: validation

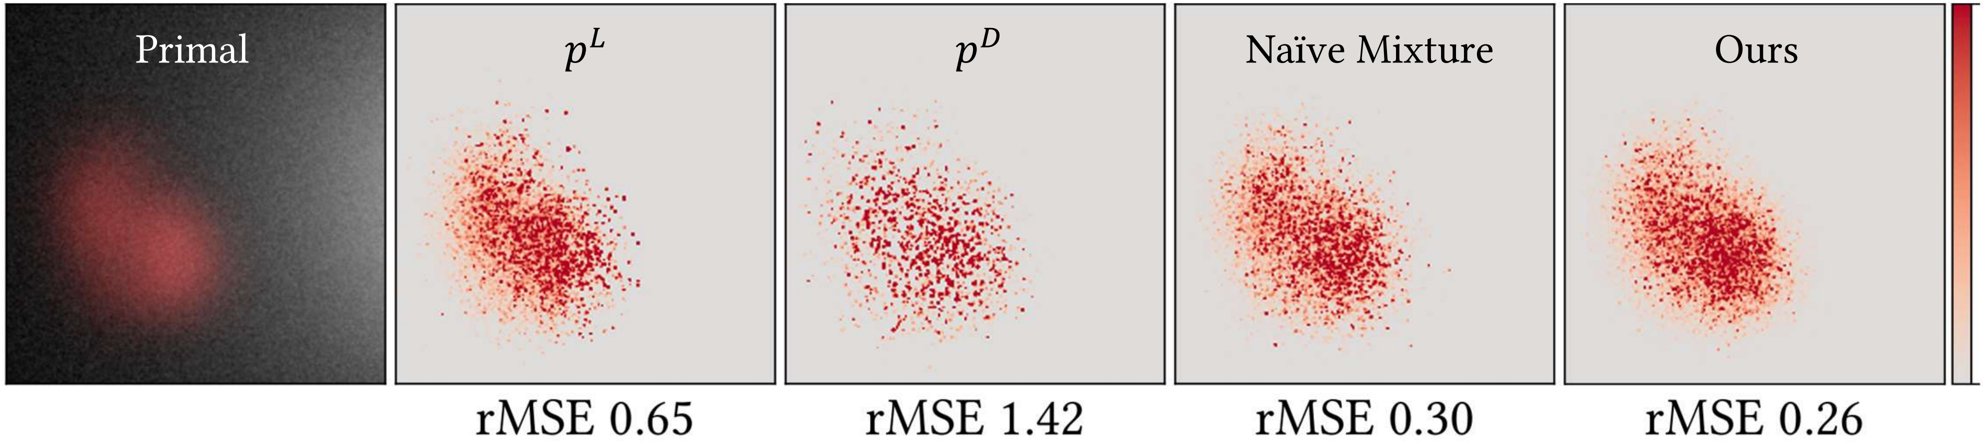

Results: gradient estimation (equal-time)

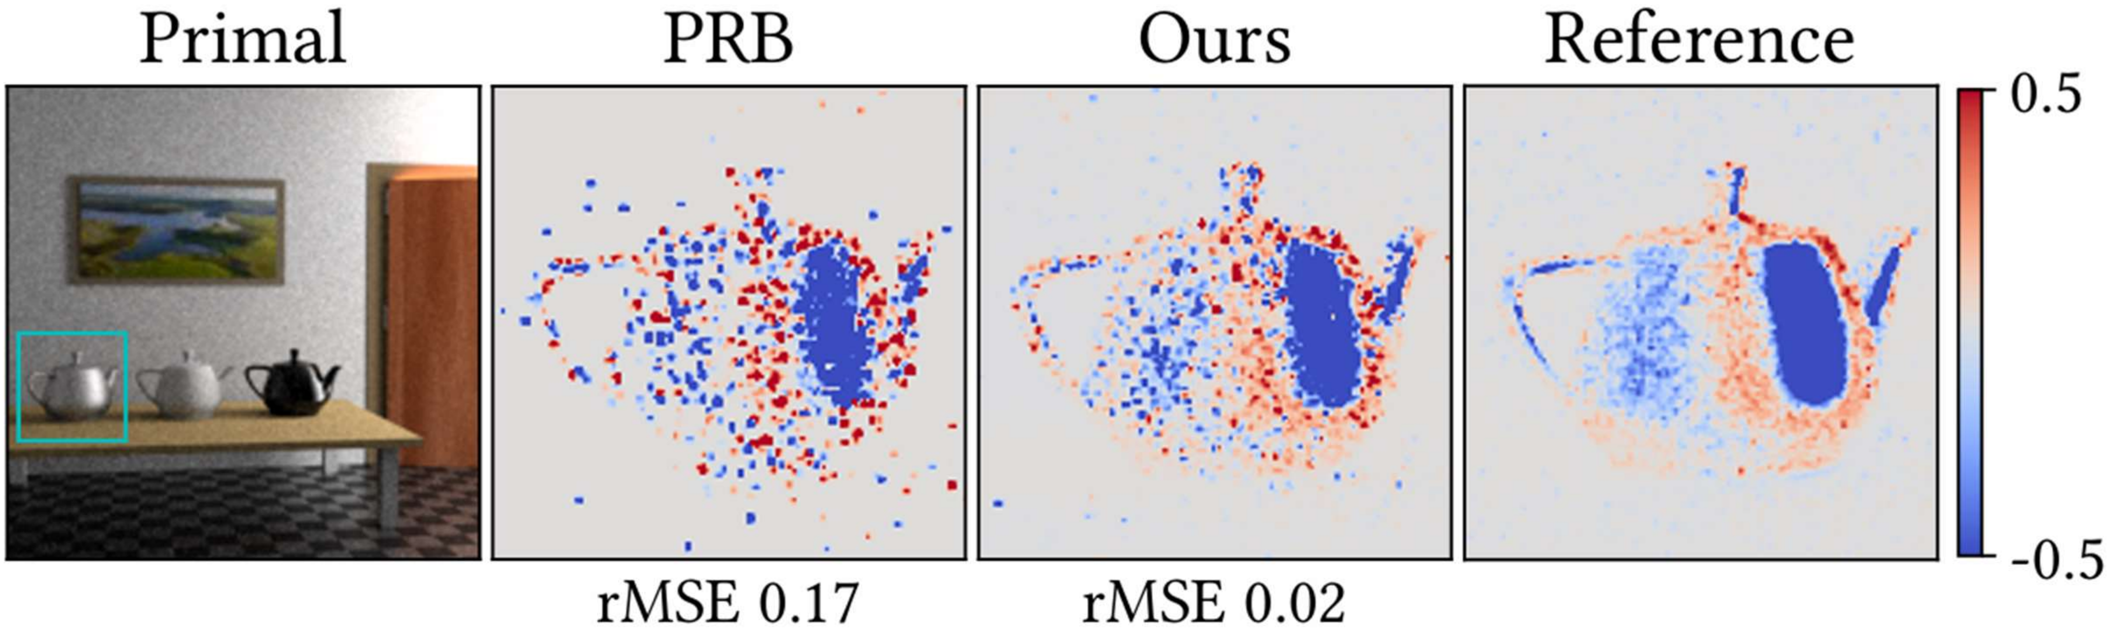

Results: inverse rendering

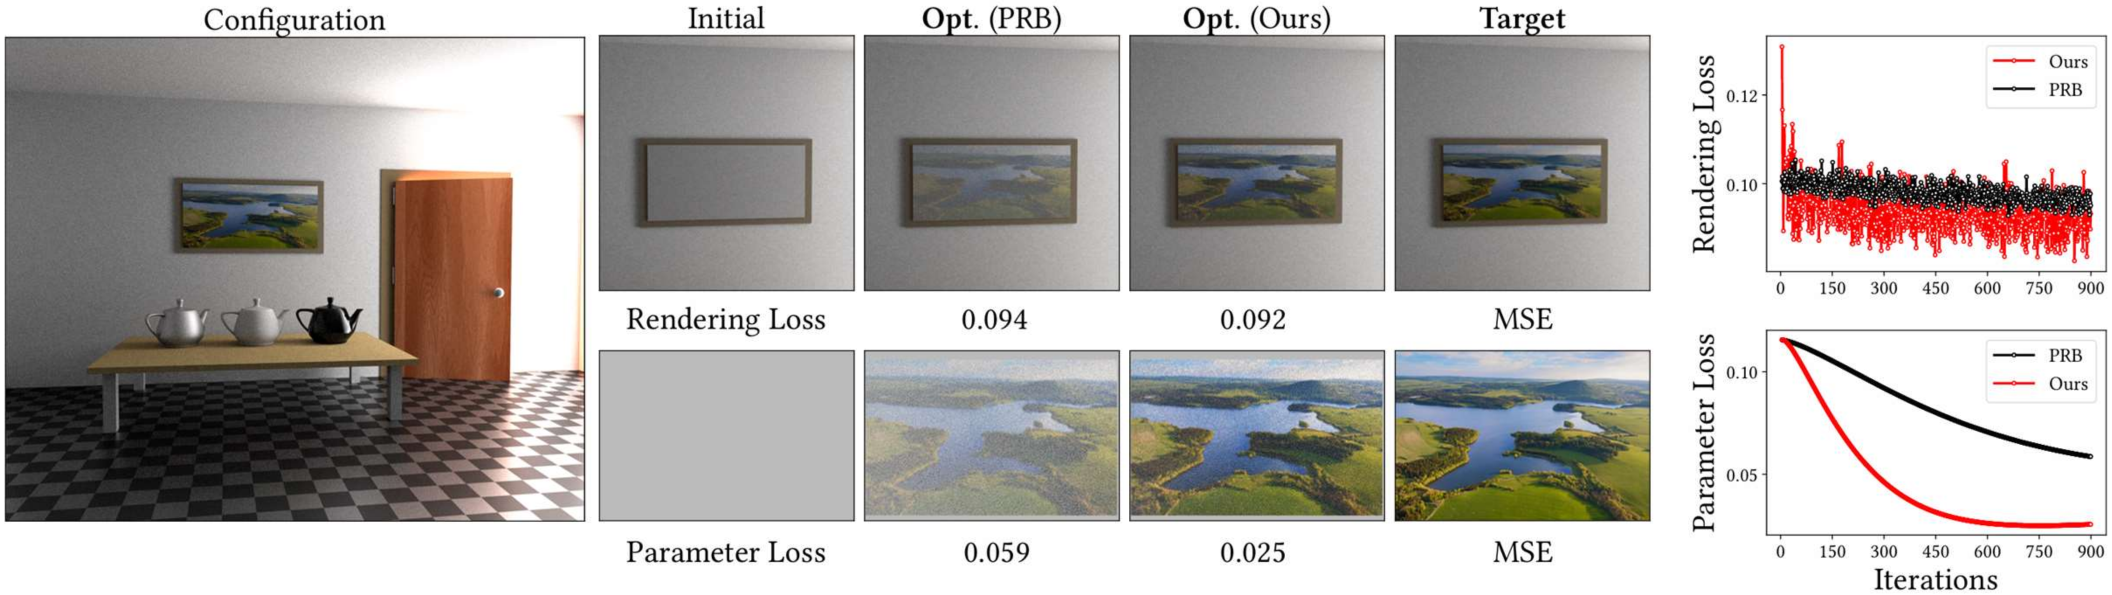

Conclusion

Importance sampling for material derivatives under global illumination

- Using p^D or p^L only has clear failure cases
- A mixture of them is more robust
- Close-form mixture weights conditioned on path prefixes
- Application in combination with path guiding

Future works

- Better guiding structures for differentiable rendering
- Better distribution sharing across iterations
- Difference between consecutive iterations
- Failure cases like pure-specular scenes

SIGGRAPH 2024

DENVER+ 28 JUL — 1 AUG

THANK YOU