SIGGRAPH ASIA 2023 SYDNEY

Extended Path Space Manifolds for Physically Based Differentiable Rendering

Jiankai Xing[1], Xuejun Hu[1], Fujun Luan[2], Ling-Qi Yan[3], Kun Xu[1]
[1]Tsinghua University & [2]Adobe & [3]UC Santa Barbara

Connecting STORIES
The 16th ACM SIGGRAPH Conference and Exhibition on Computer Graphics and Interactive Techniques in Asia
CONFERENCE 12 - 15 December 2023
EXHIBITION 13 - 15 December 2023
ICC, Sydney, Australia
ASIA.SIGGRAPH.ORG/2023

Differentiable Rendering

Forward Rendering

$$y = f(x)$$

$$\frac{\partial y}{\partial x} = ?$$

Differentiable Rendering

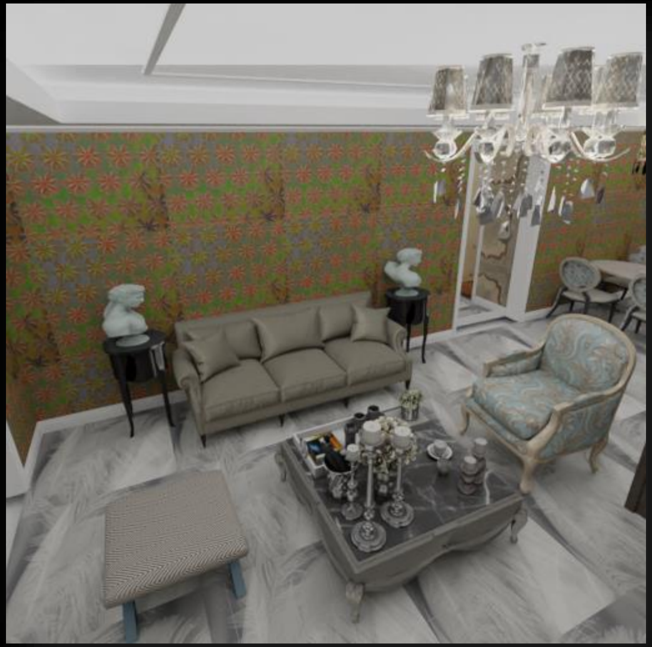

x: Geometry Material Light Camera ……

y: Rendered Image

Application: Inverse Rendering & 3D AIGC

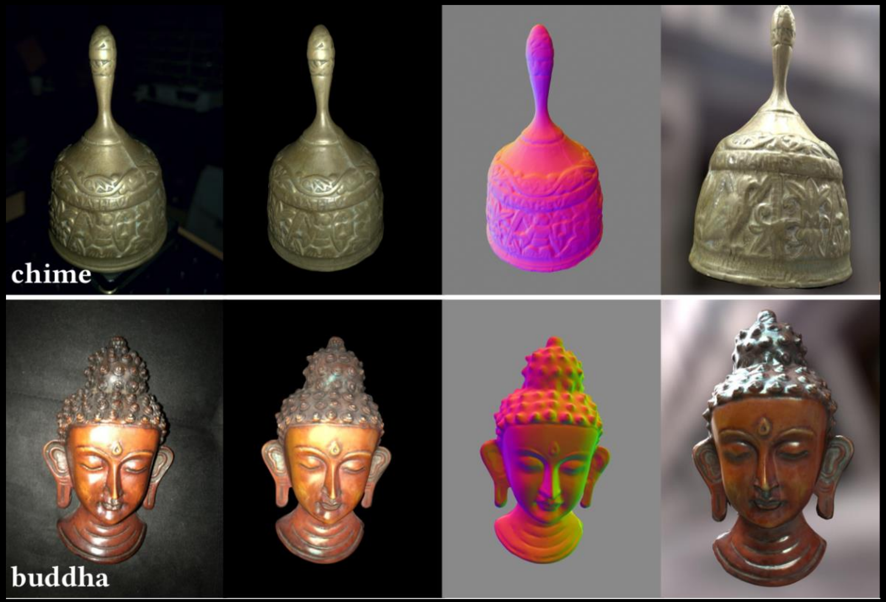

Inverse Rendering

Recover scene representation from images/videos through analysis by synthesis

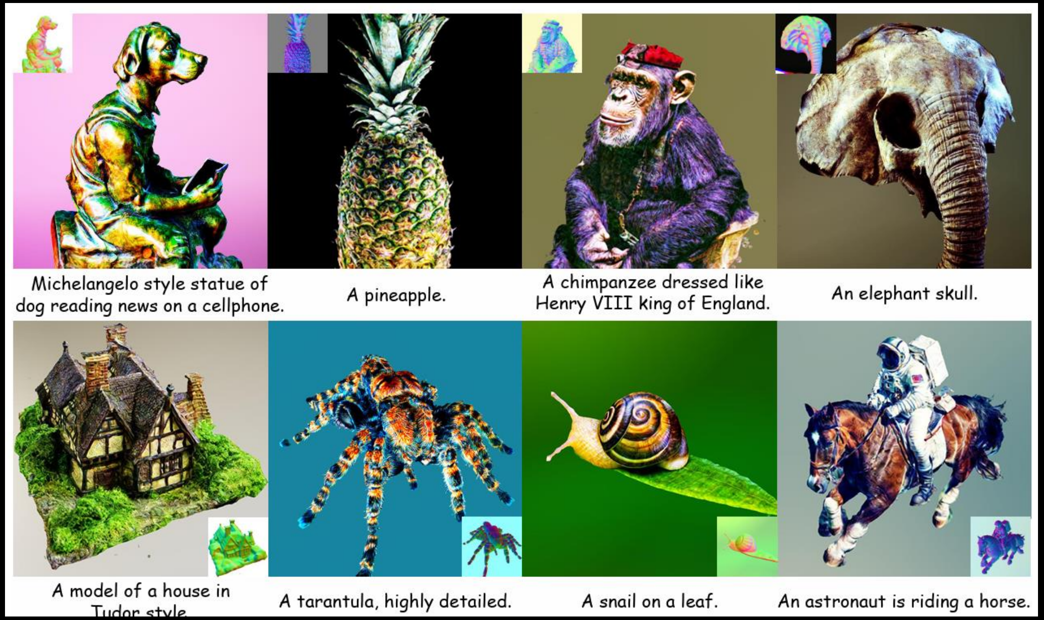

3D AIGC

Optimize 3D representation using 2D images priors Through SDS/VSD Loss

Differentiable Rendering Pipeline

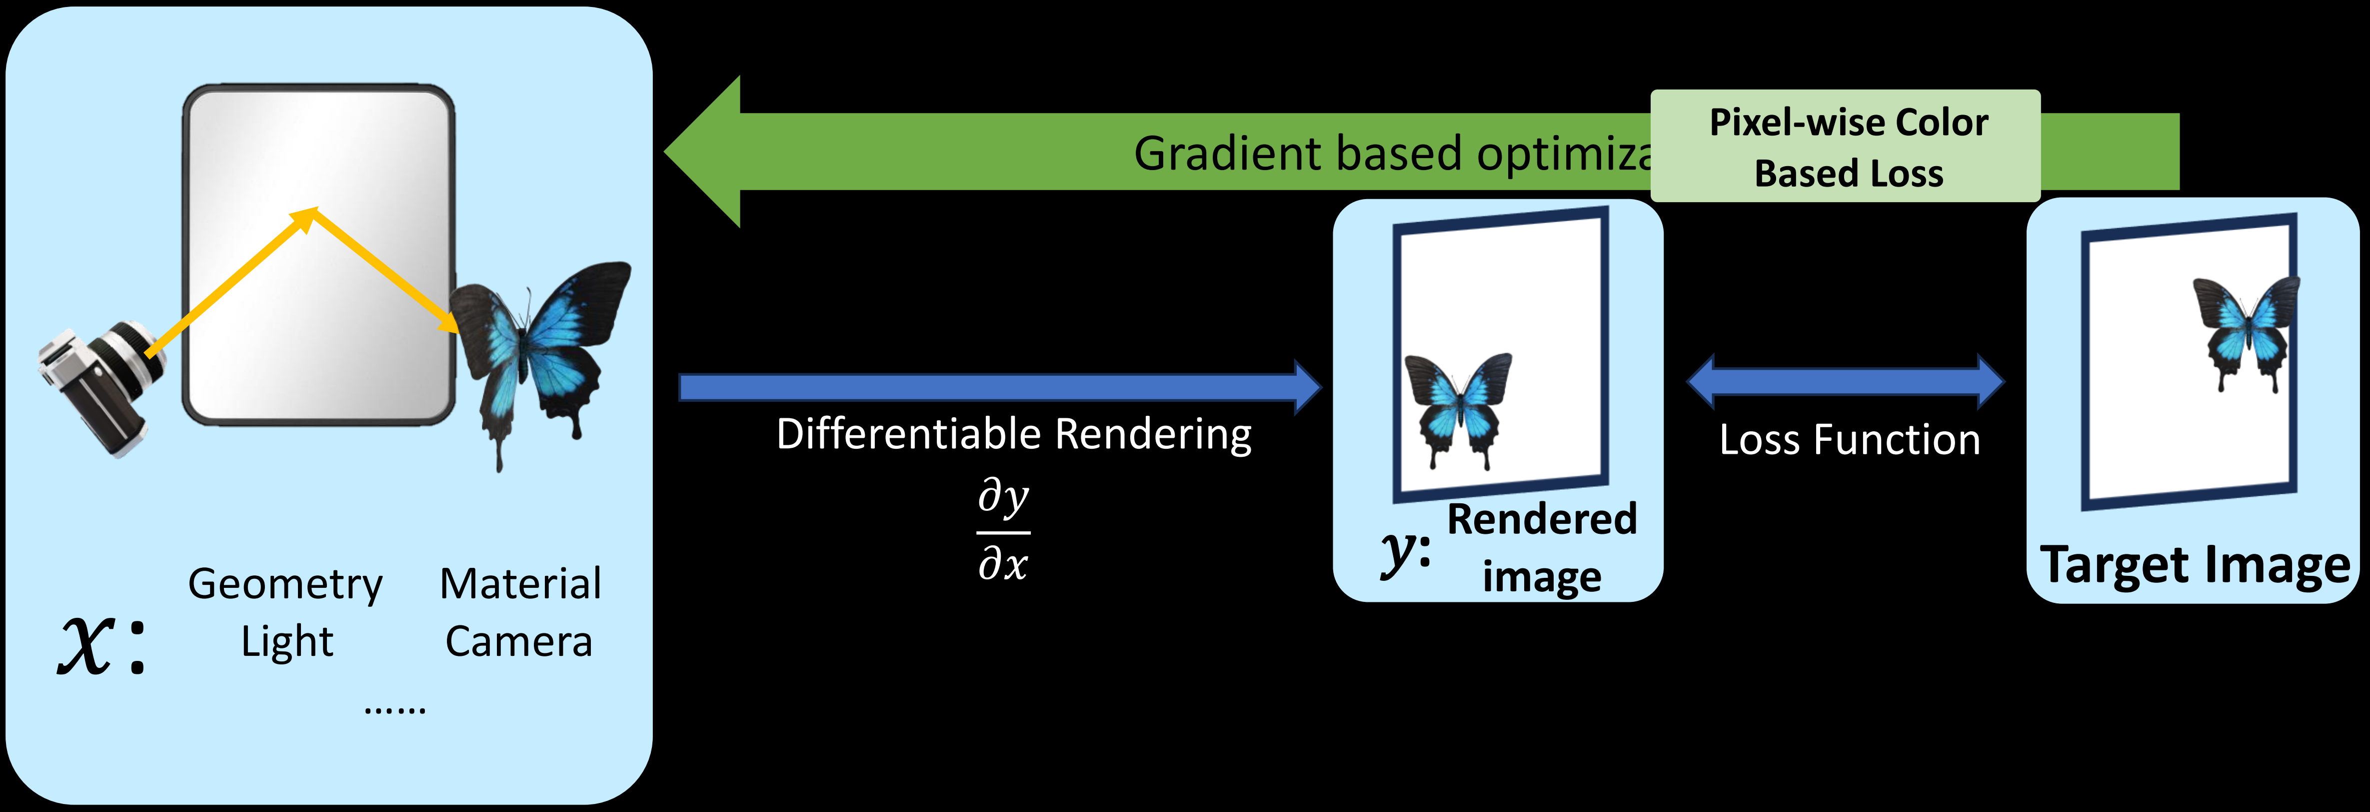

Differentiable Rendering Pipeline: Rasterization

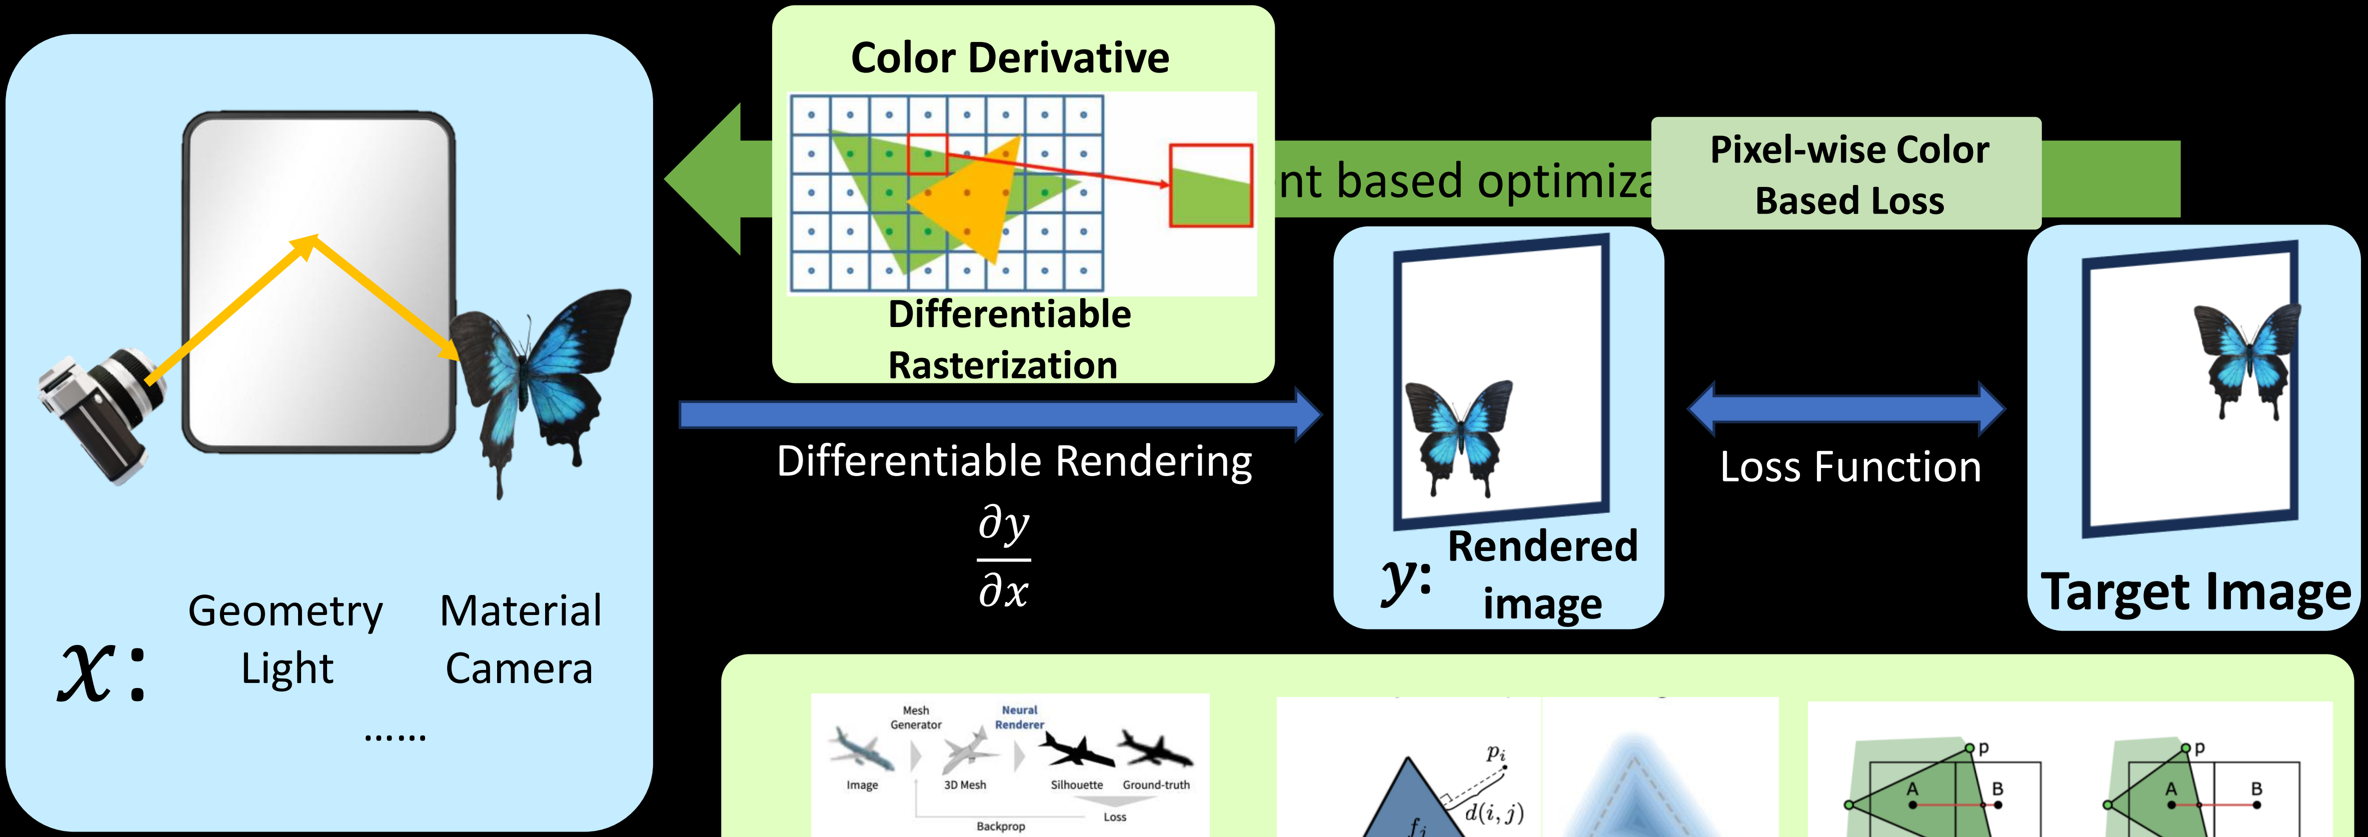

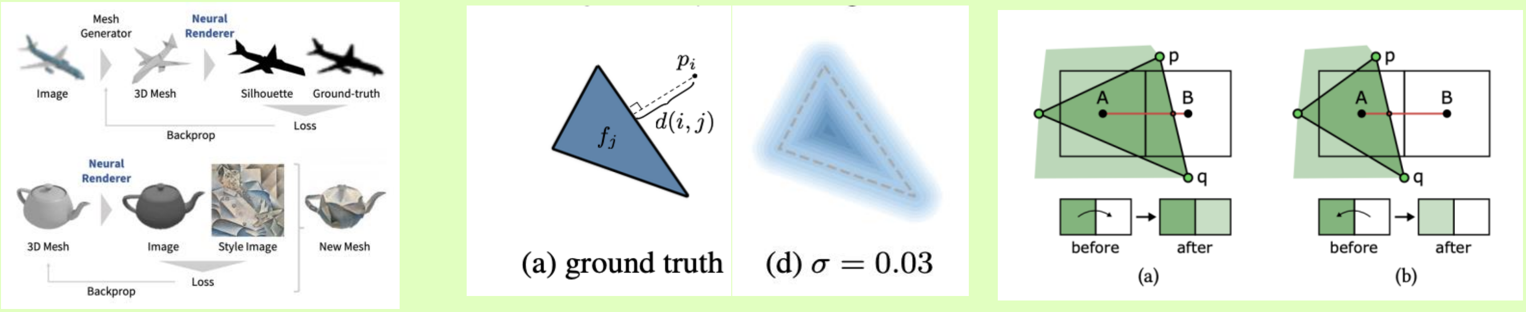

[Kato 18] N3MR
Approximate Gradient

Soft Rasterizer, PyTorch3D
Approximate Rendering

Nvdiffrast
Analytical Antialiasing

Differentiable Rendering Pipeline: PBR/PT

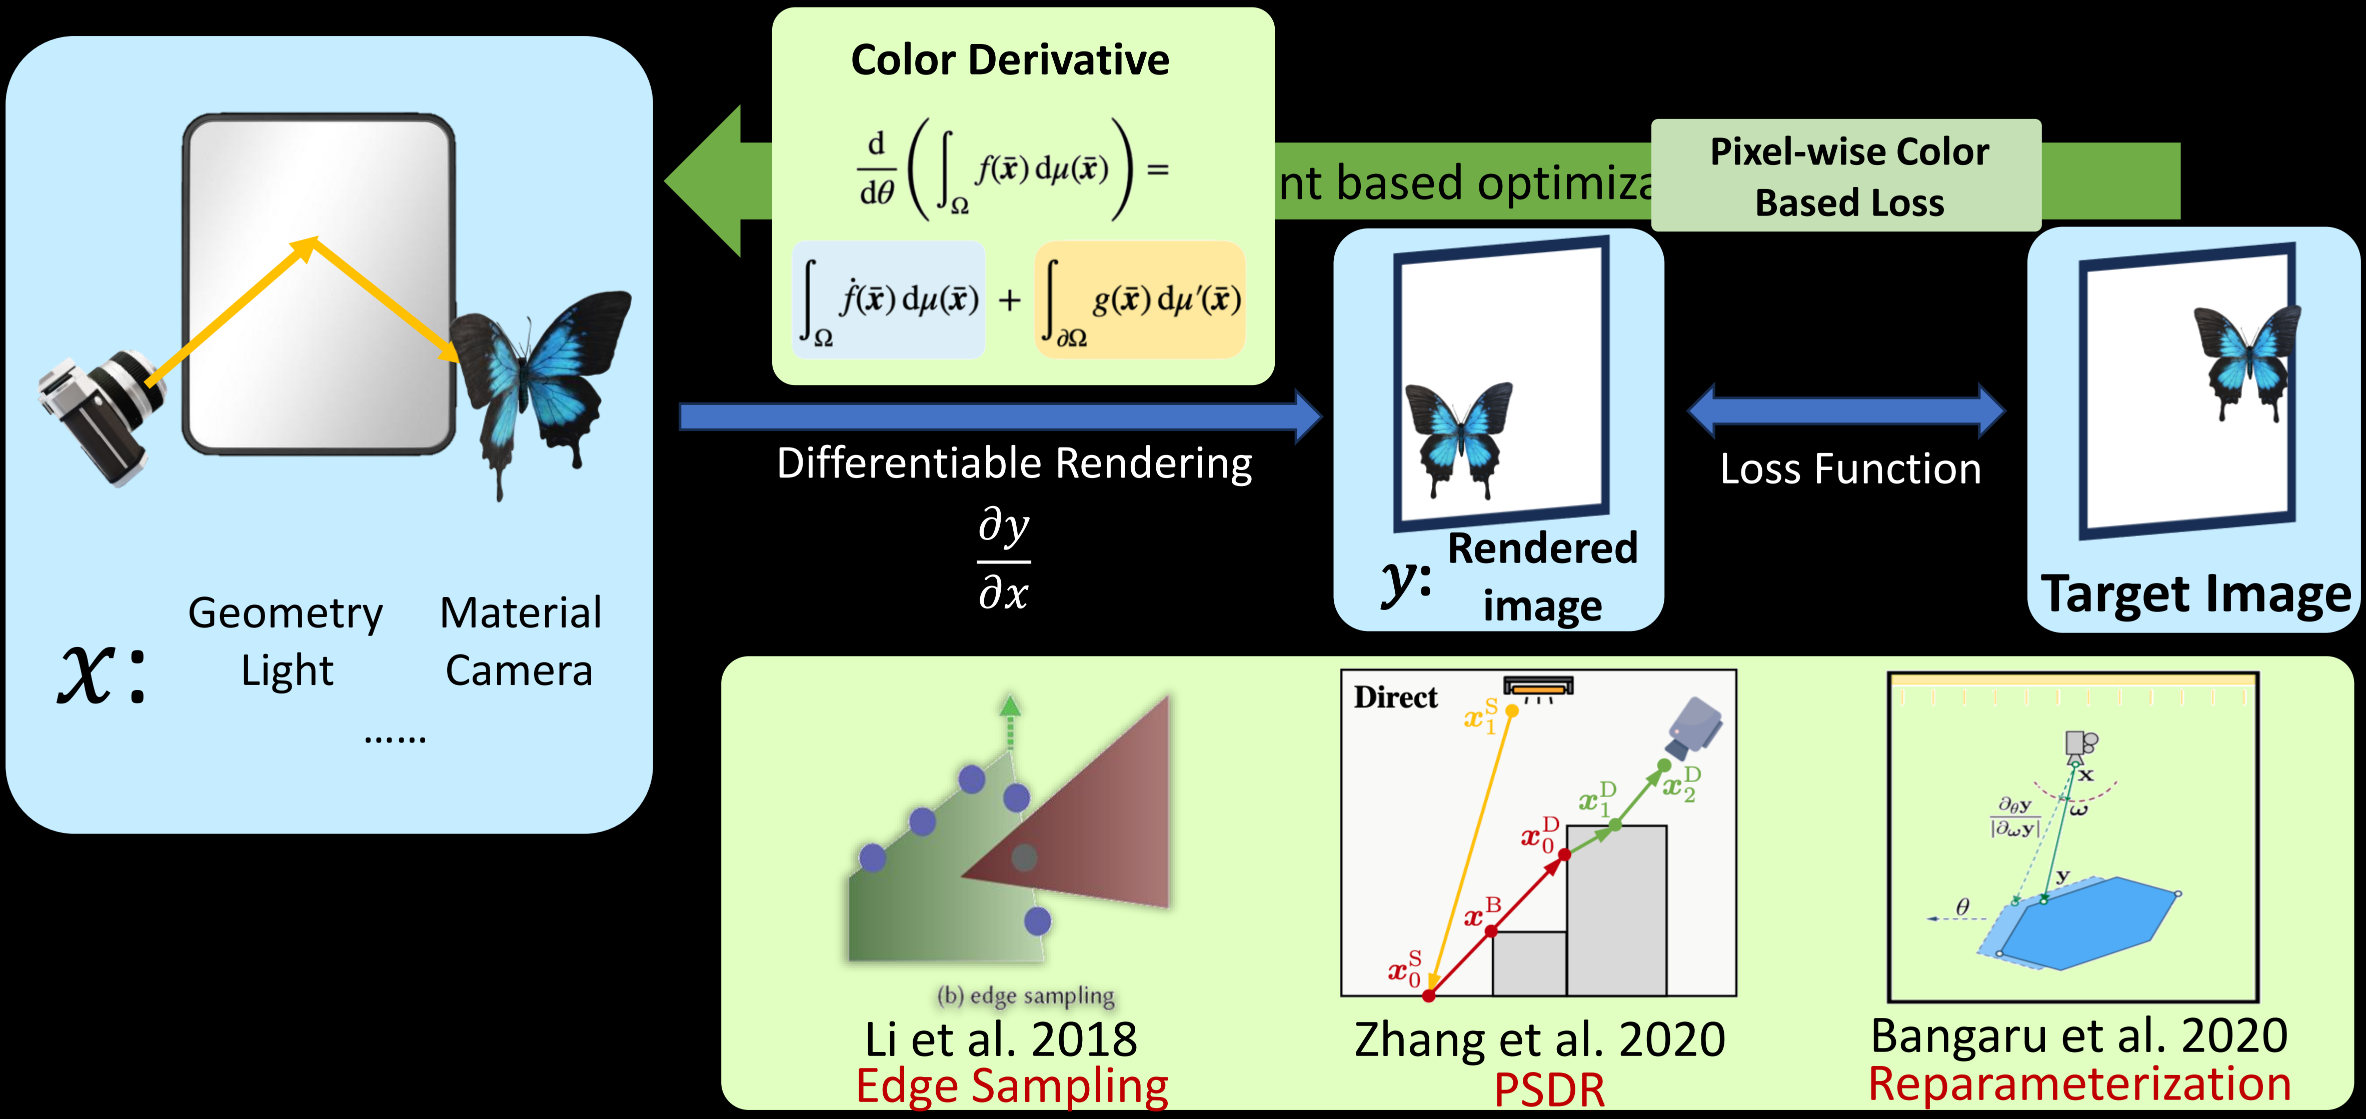

Problems: Loss Locality & Gradient Sparsity

Rendered Image

Target Image

Loss Locality

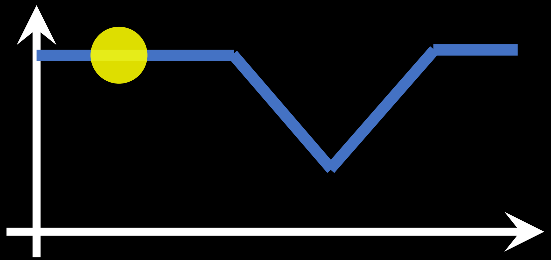

Overall L2 Loss with
Object Translation

Gradient Sparsity

$\theta: x - axis$ translation

Gradient

Some methods may help
- multi-resolution
- plateau reduced
  [Fischer 2022]
- large-steps
  [Nicolet 2021]

But not enough!

A Novel Pipeline

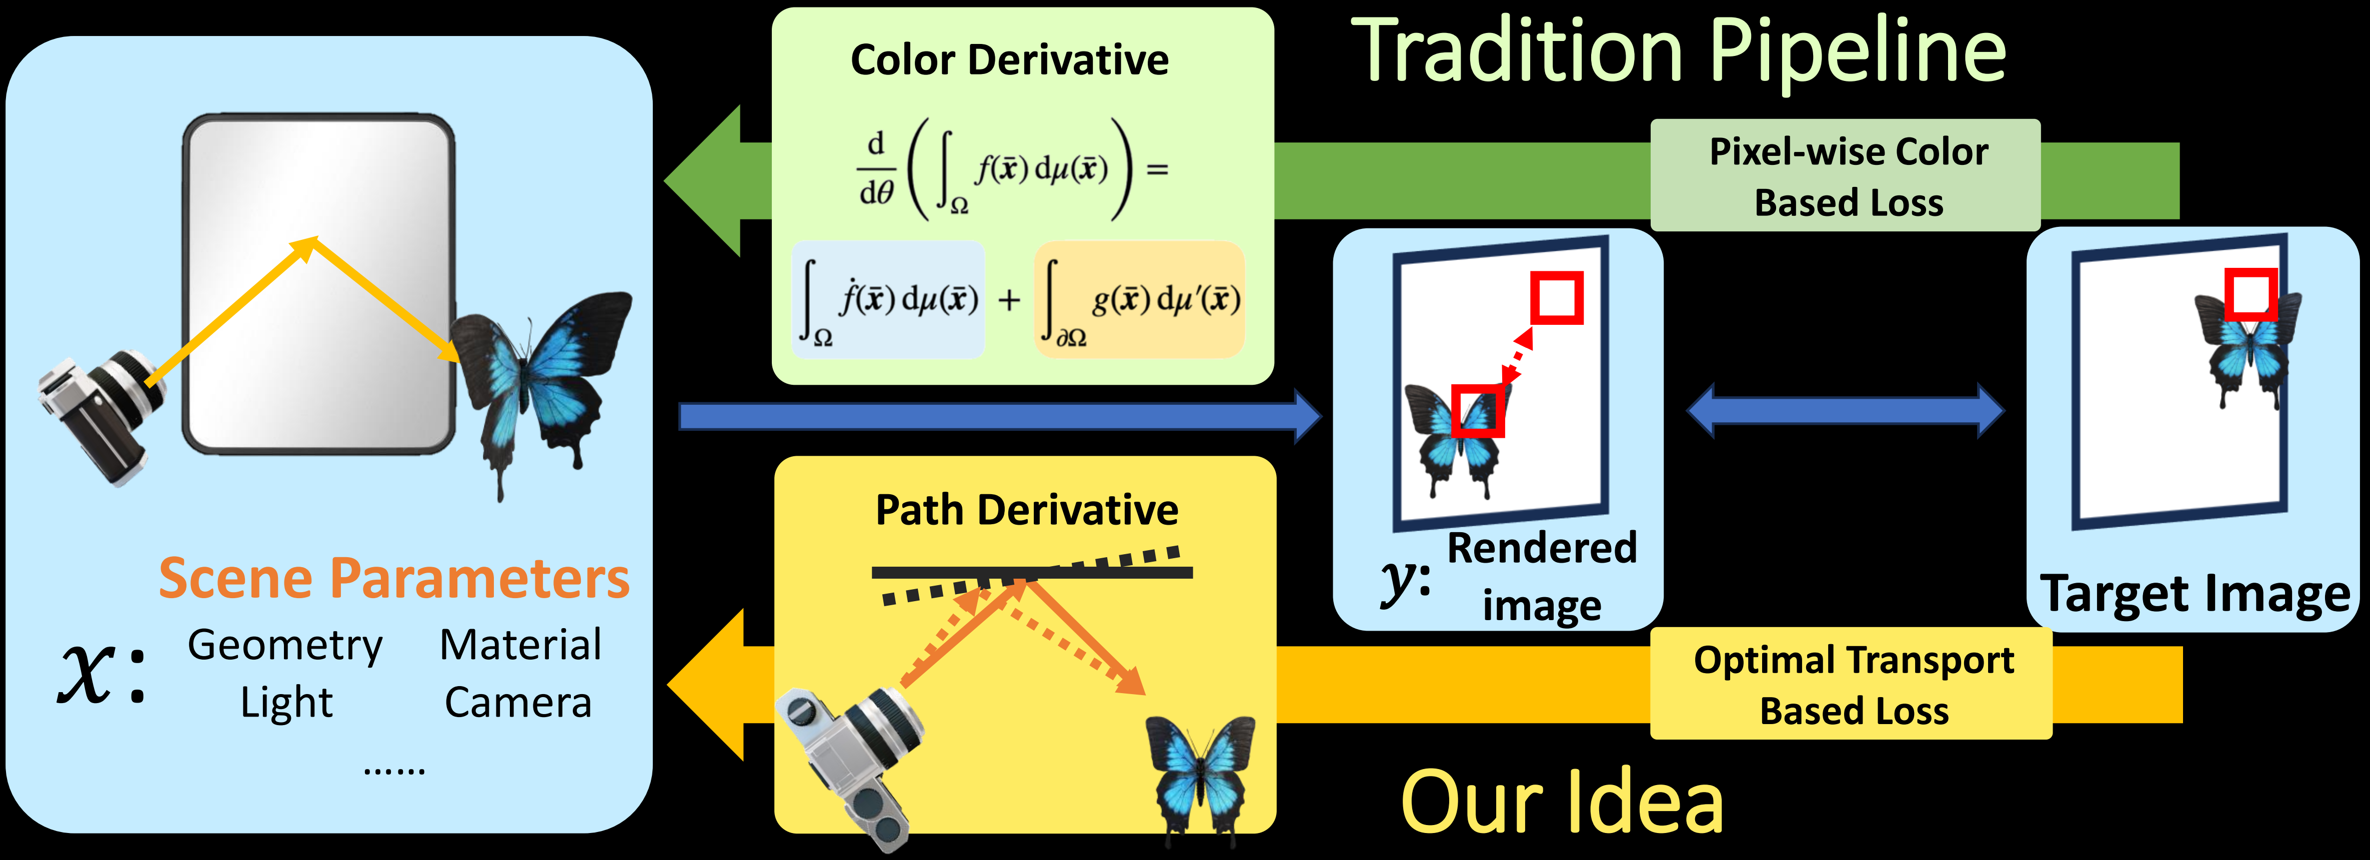

Optimal Transport Based Loss Function [Xing 2022]

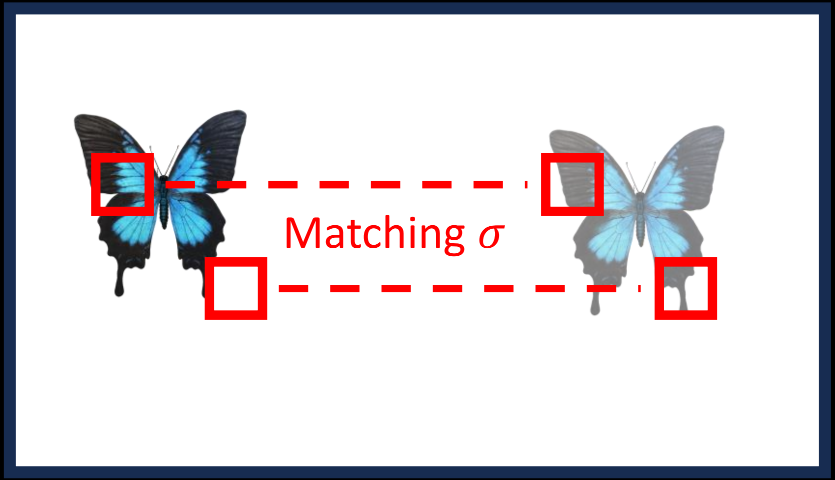

Current & Target Image

$$\min_{\sigma \in Perm(n)} \frac{1}{n} \sum_i d\left(\boldsymbol{P}_i^{\text{render}}, \boldsymbol{P}_{\sigma(i)}^{\text{target}}\right)$$

$$d\left(\boldsymbol{P}_i, \boldsymbol{P}_j\right) = \left\|\boldsymbol{P}_i - \boldsymbol{P}_j\right\|_2^2$$

Solved by optimal transport

$$\boldsymbol{P}_i = \begin{bmatrix} r \\ g \\ b \\ x \\ y \end{bmatrix}$$

Color $\boldsymbol{c}$ (r, g, b)

Screen Space Position $\boldsymbol{p}$ (x, y)

$$\mathrm{L}^{our} = \frac{1}{n} \sum \left\| \boldsymbol{P}_i^{\text{render}} - \boldsymbol{P}_{\sigma(i)}^{\text{target}} \right\|^2$$

Traditional Loss vs Our Loss

$$\boldsymbol{P}_i = \begin{bmatrix} r \\ g \\ b \\ \lambda x \\ \lambda y \end{bmatrix}$$

Color $\boldsymbol{c}$

Screen Space Position $\boldsymbol{p}$

Traditional Loss

$$\mathrm{L} = \frac{1}{n}\sum\left\|\mathbf{I}_i^{\mathrm{render}} - \mathbf{I}_i^{\mathrm{target}}\right\|^2 = \frac{1}{n}\sum\left\|\boldsymbol{P}_i^{\mathrm{render}} - \boldsymbol{P}_i^{\mathrm{target}}\right\|^2$$

Our Loss

$$\mathrm{L}^{our} = \frac{1}{n}\sum\left\|\boldsymbol{P}_i^{\mathrm{render}} - \boldsymbol{P}_{\sigma(i)}^{\mathrm{target}}\right\|^2$$ from a matching

Traditional loss is equivalent to our loss with identity matching

Backpropagate to Scene Parameter $\boldsymbol{\theta}$

$$\boldsymbol{P}_i = \begin{bmatrix} r \\ g \\ b \\ x \\ y \end{bmatrix}$$

Color $\boldsymbol{c}$

Screen Space Position $\boldsymbol{p}$

$$\mathrm{L}^{our} = \frac{1}{n} \sum \left\| \boldsymbol{P}_i^{\text{render}} - \boldsymbol{P}_{\sigma(i)}^{\text{target}} \right\|^2$$

$$\frac{\partial \mathrm{L}^{our}}{\partial \theta} = \frac{\partial \mathrm{L}^{our}}{\partial \boldsymbol{c}} \boxed{\frac{\partial \boldsymbol{c}}{\partial \theta}} + \frac{\partial \mathrm{L}^{our}}{\partial \boldsymbol{p}} \boxed{\frac{\partial \boldsymbol{p}}{\partial \theta}}$$

Color Derivatives (Follow traditional work)

How to compute position derivatives with respect to scene parameter θ?

Path Derivatives

$$\boldsymbol{p} \rightarrow \boldsymbol{x} \rightarrow \boldsymbol{\theta}$$

Pixel Position **Path** **Scene Parameters**

$$\frac{\partial \boldsymbol{p}}{\partial \boldsymbol{\theta}} = \frac{\partial \boldsymbol{p}(\boldsymbol{x})}{\partial \boldsymbol{x}} \frac{\partial \boldsymbol{x}}{\partial \boldsymbol{\theta}}$$

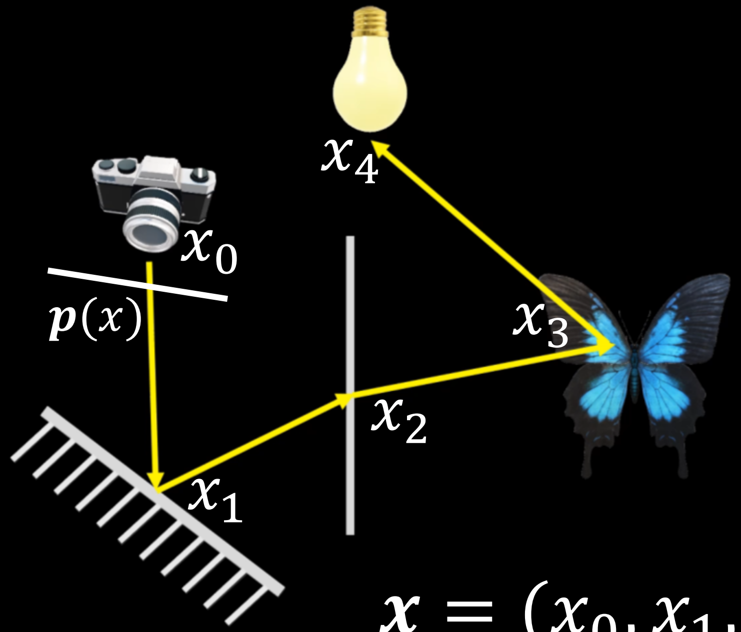

$$\boldsymbol{x} = (x_0, x_1, x_2, x_3, x_4)$$

Regard pixel position $\boldsymbol{p}$ as the path $\overline{\boldsymbol{x}}$ interaction with image plane

Path Derivatives: When we change scene parameter $\boldsymbol{\theta}$, how will the path $\boldsymbol{x}$ be changed

Simple Path Derivatives [Xing 2022]

Path $\boldsymbol{x} = (x_0, x_1)$, where x_1 is the primary intersection

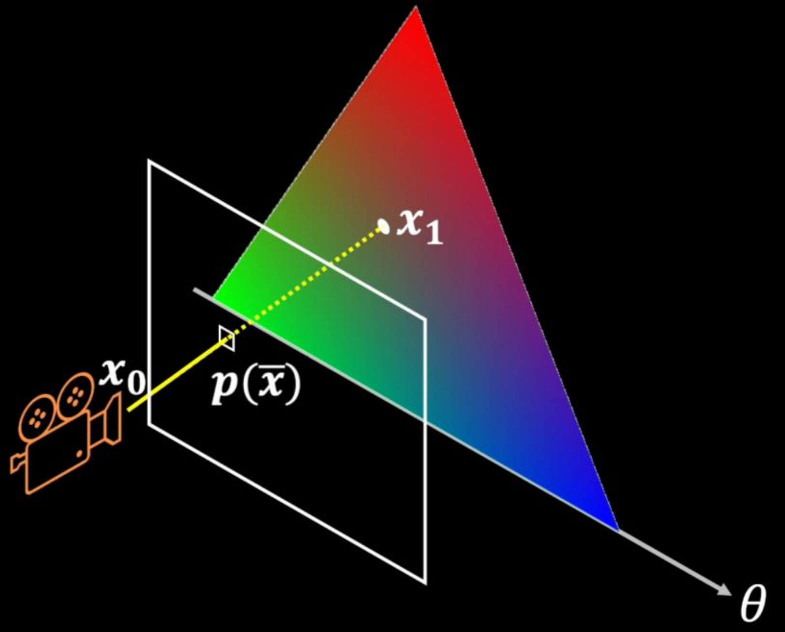

When changing scene parameter $\boldsymbol{\theta}$,

Stick x_1 to the underlying geometry

Cannot handle longer light paths involving global illumination effects!

Complex Path in PT

- In path tracing, paths are composed of multiple vertex and generated by various sampling strategy
- No determined relationship between scene parameters and path
  - Unable to compute derivatives of path with respect to parameters.

Our Main Task:

Build a determined relationship between scene parameters and complex light path.

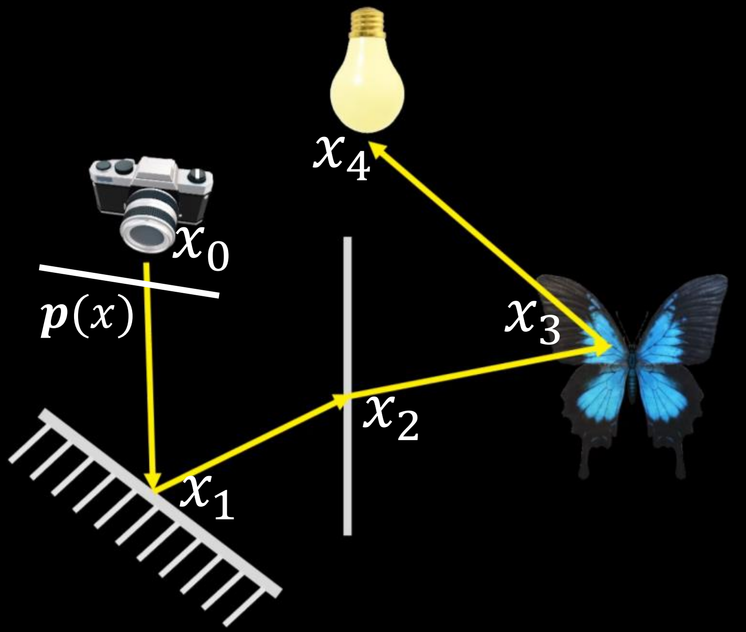

$$\boldsymbol{x} = (\boldsymbol{x}_0, \boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3, \boldsymbol{x}_4)$$

$\boldsymbol{\theta}$: mirror rotation

Inspiration: Path Space Manifold [Jakob 2012]

- Valid paths usually lie in low dimension manifolds.
- The specular middle points can be **uniquely** determined by the diffuse endpoints.

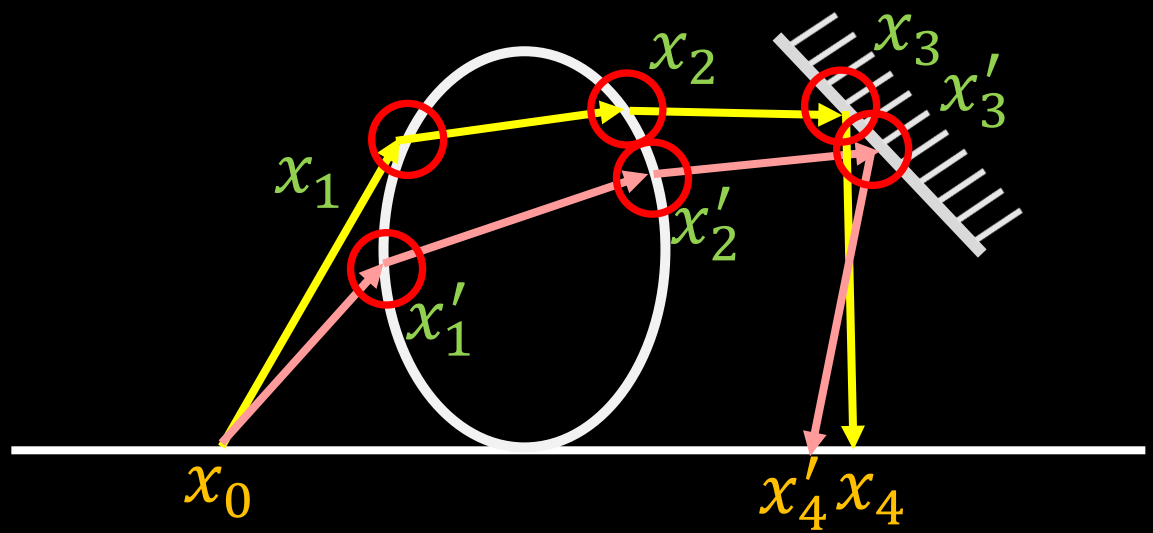

$$\boldsymbol{x} = (x_0, x_1, x_2, x_3, x_4)$$

Constraints $C(\boldsymbol{x}) = 0$

Implicit Function Theorem

$$\frac{\partial x_i}{\partial x_j}$$

Extended Path Space Manifold (EPSM)

Path Space Manifold

$$C(\boldsymbol{x}) = 0, \frac{\partial x_i}{\partial x_j}$$

Efficient Sampling

Forward Rendering

Extended PSM

$$C(\boldsymbol{x}, \theta) = 0, \frac{\partial \boldsymbol{x}}{\partial \theta}$$

Path Derivatives

Differentiable Rendering

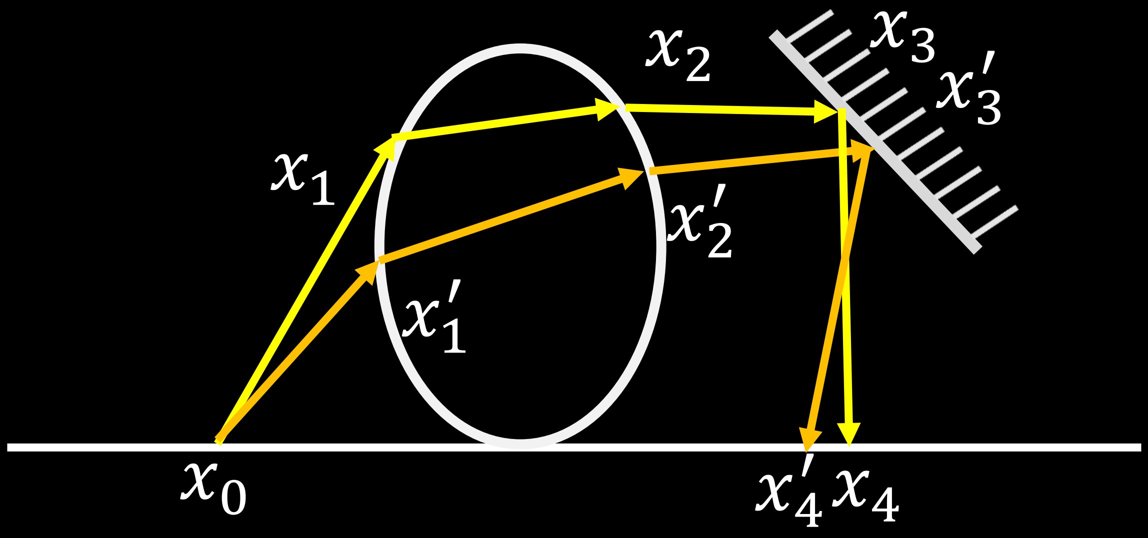

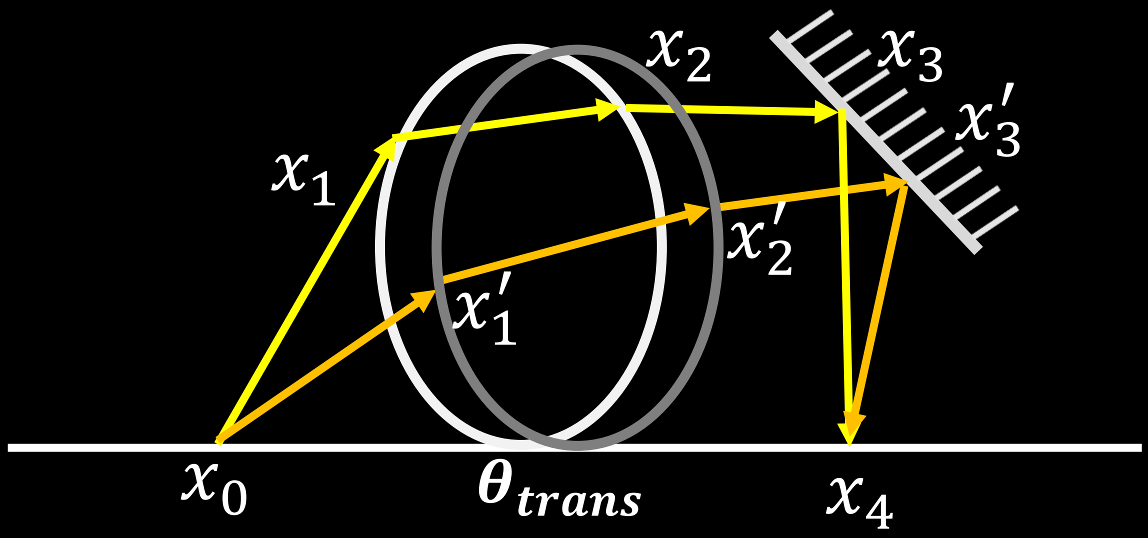

How to build $C(x,\theta)$=0

Constraints $C(x,\theta)$:
Represents how we want
the path to be changed

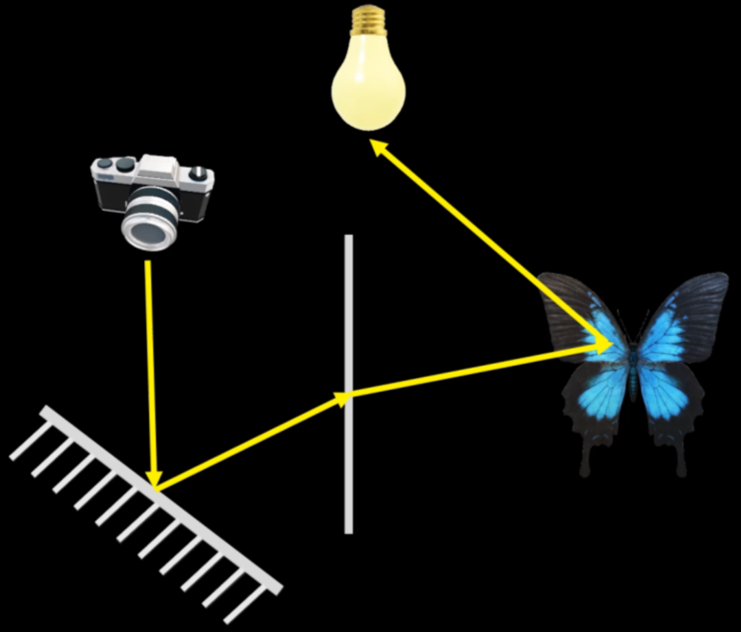

Goal of Building Constraint

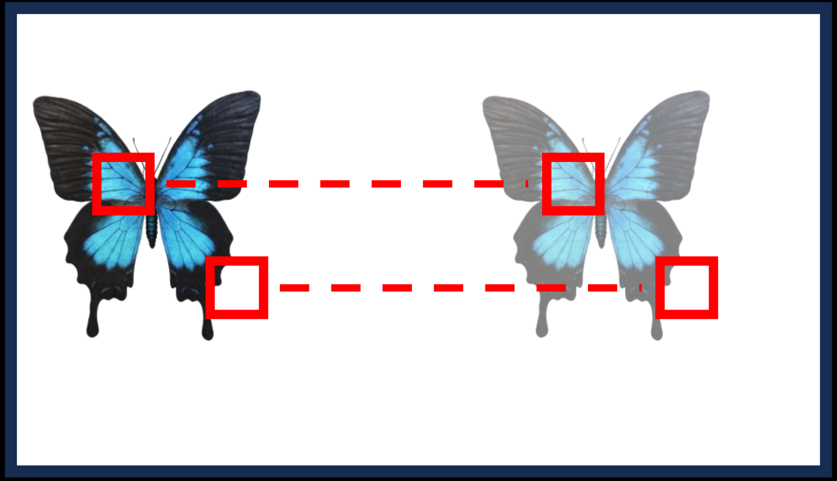

Move path geometry
Keep radiance

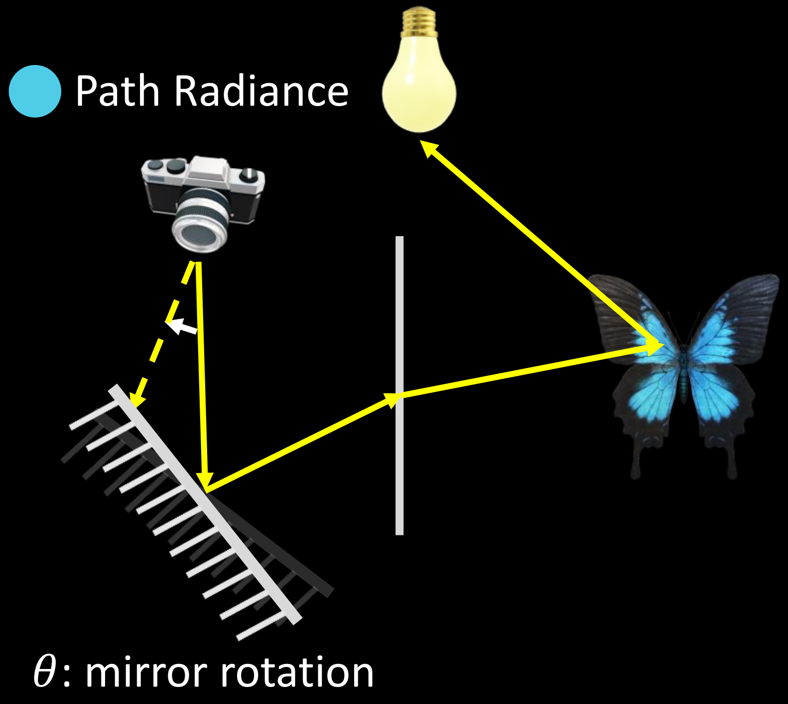

Goal of Building Constraint

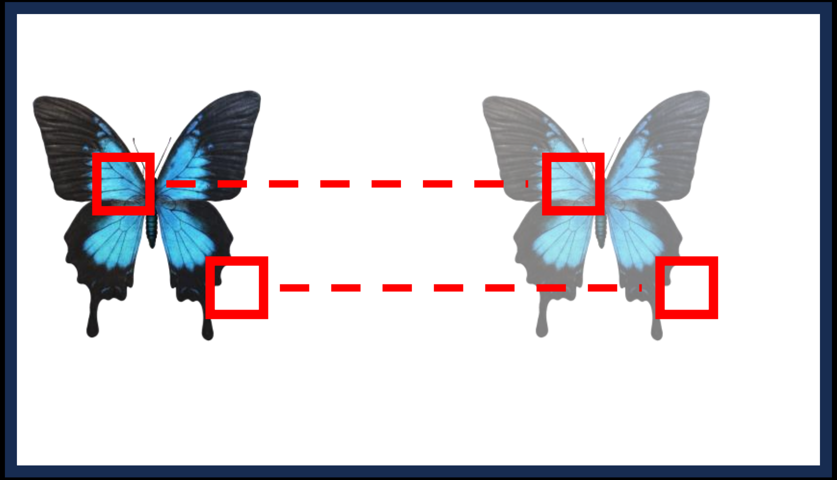

Move path geometry

Keep path effect

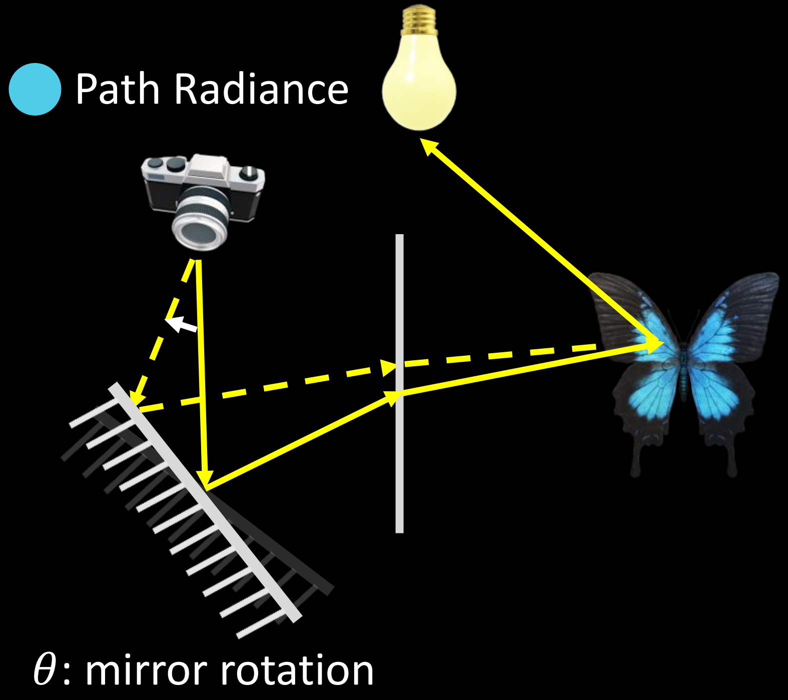

Example: Keep Path Effect

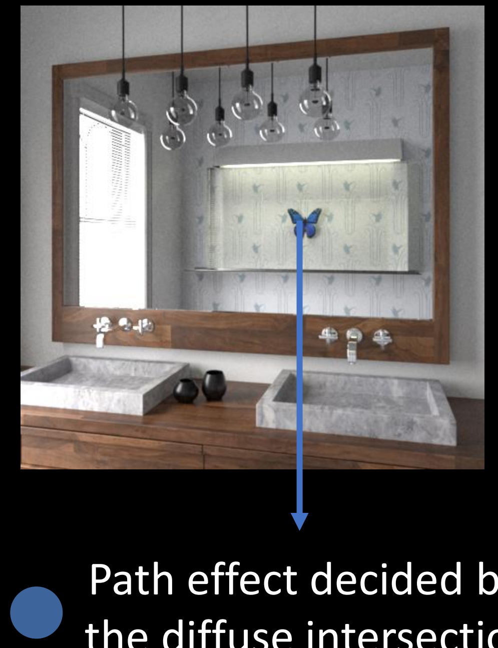

Path effect decided by the diffuse intersection

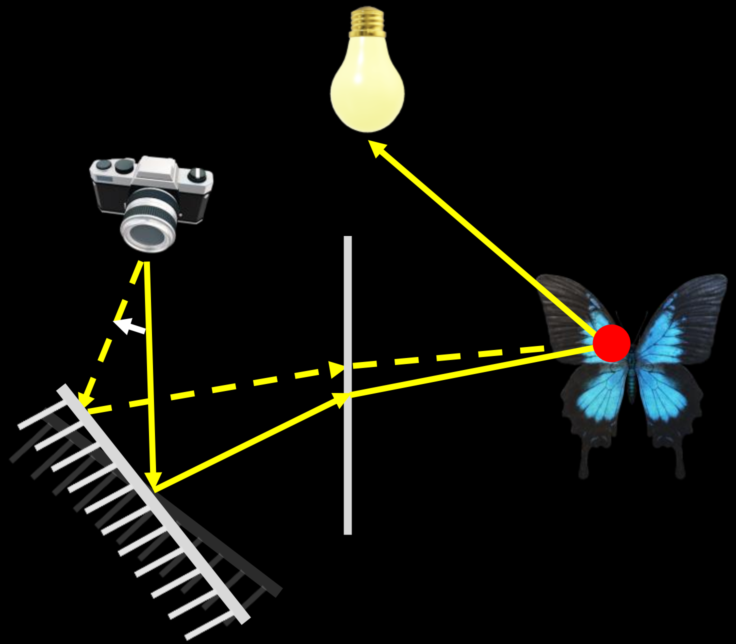

θ: mirror rotation

Fix Position Constraint

Fix Position Constraint

Visualization

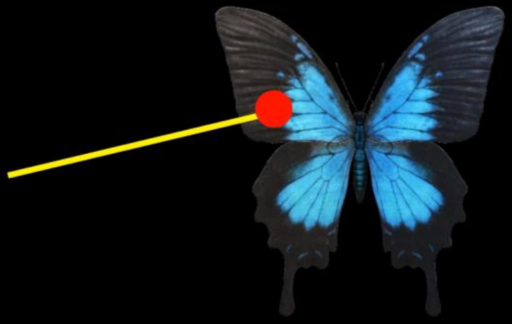

Mathematics

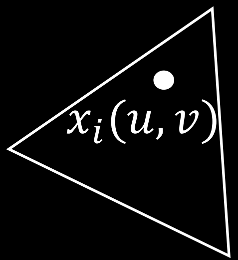

$u = Const$
$v = Const$

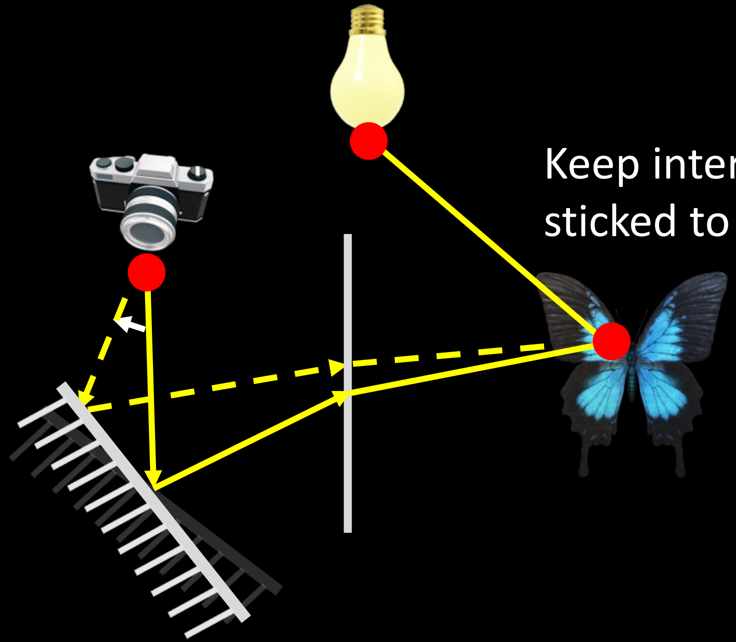

Half Vector Constraint

Fix Position Constraint

Half Vector Constraint

Visualization

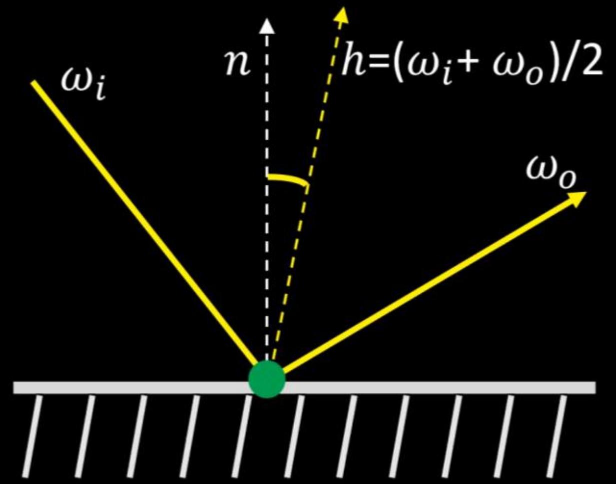

Mathematics

$$h=(\omega_i+\omega_o)/2$$
$$< n, h > = Const$$

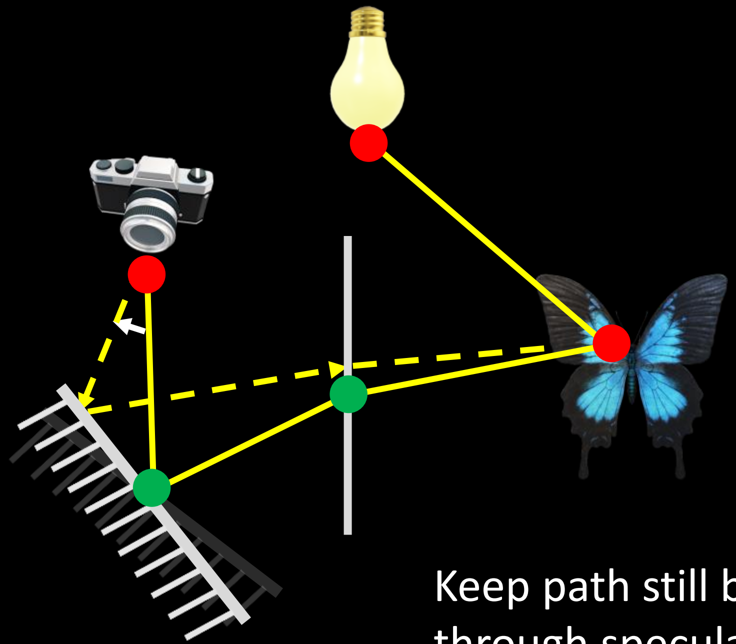

θ: mirror rotation

Keep path still be valid through specular /transparent /glossy intersection

Example: Path Moving

- Fix Position Constraint
- Half Vector Constraint

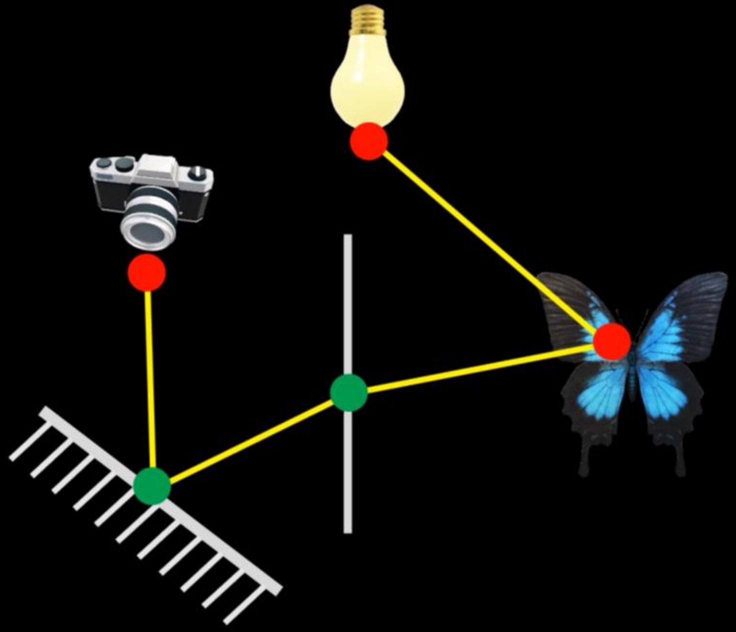

θ: mirror rotation

Derivatives Computing

- For length n path $\boldsymbol{x} = (x_1, x_2, \dots, x_n)$ (all surface intersections)
- We enforce n 2D Constraint on it:

$$C(\boldsymbol{x}, \theta) = \{C_i(\boldsymbol{x}, \theta) = \mathbf{0} \quad \boldsymbol{i} = \mathbf{1} \dots \boldsymbol{n}\}$$

- Given θ, we can uniquely determine a path $\boldsymbol{x}$
- By implicit function theorem:

$$\frac{\partial \boldsymbol{x}}{\partial \theta} = -\left(\frac{\partial C}{\partial \boldsymbol{x}}\right)^{-1} \frac{\partial C}{\partial \theta}$$

Path Types and Constraints in EPSM

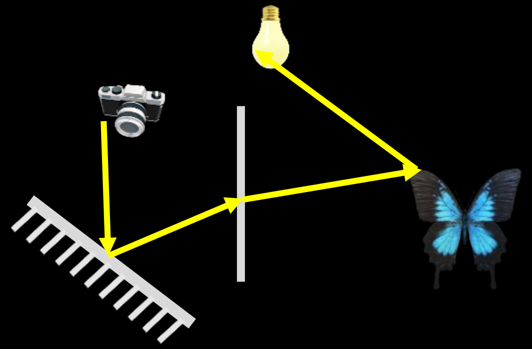

General Path

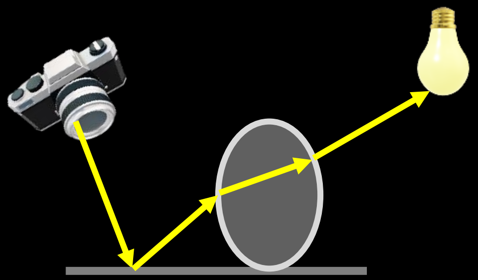

Caustic Path

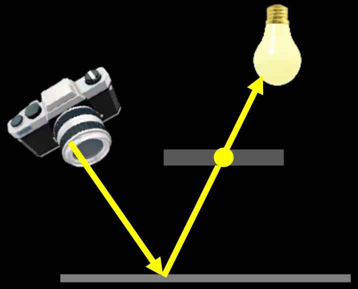

Shadow Path

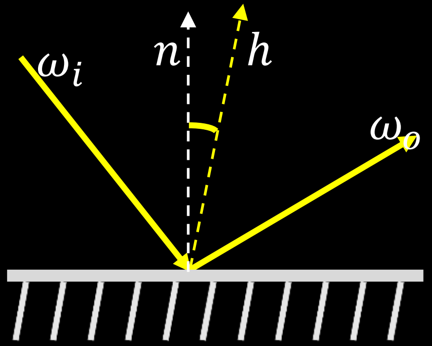

Half Vector Constraint

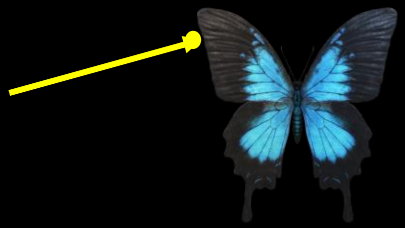

Fix Position Constraint

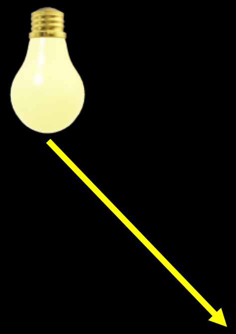

Fix Direction Constraint

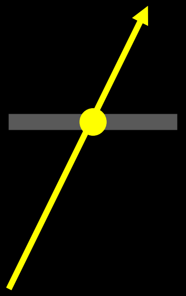

Colinear Constraint

Results

For complex path

Bathroom

Path Type: General Path
Optimize parameter:
8 Object Translation

Init

Target

PRB[Vicini 2021]

PRDPT[Fischer 2022]

Ours

Highlight

Path Type: General Path
Optimize parameter:
- Glass Slab Rotation
- Light Translation

Init

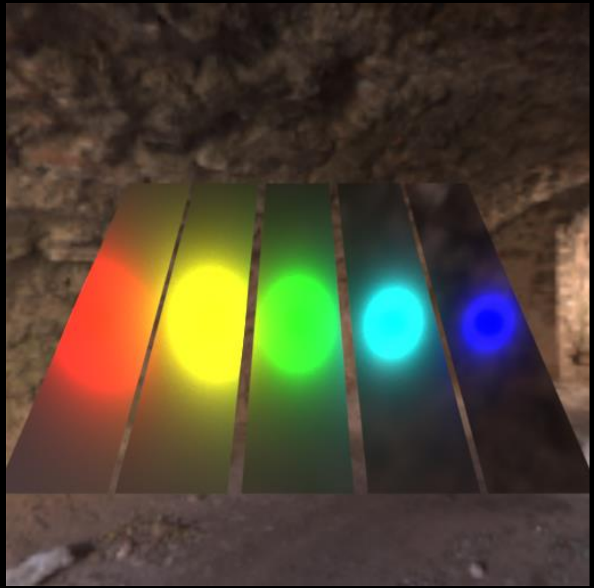

Target

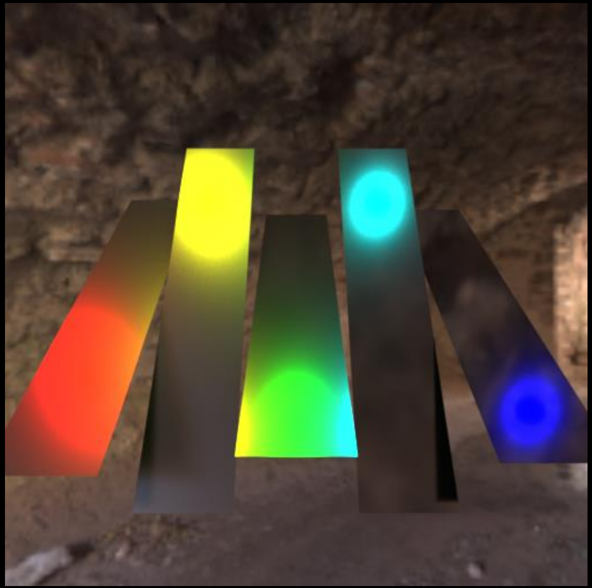

PRB[Vicini 2021]

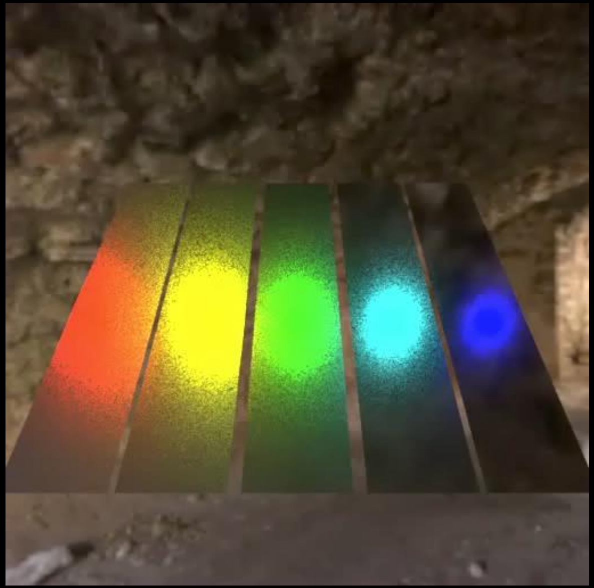

PRDPT[Fischer 2022]

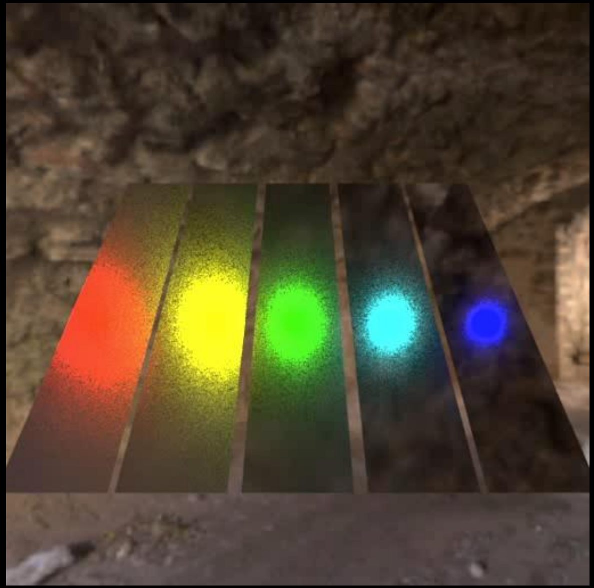

Ours

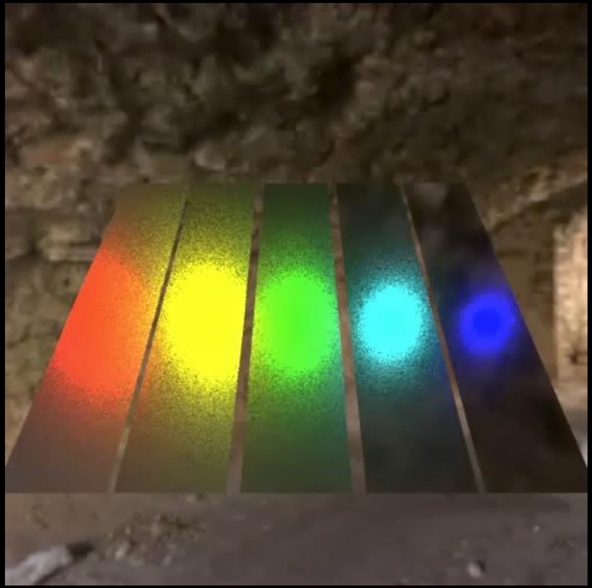

CornellBox

Path Type: Caustic Path
Optimize parameter:
Light Position

Init

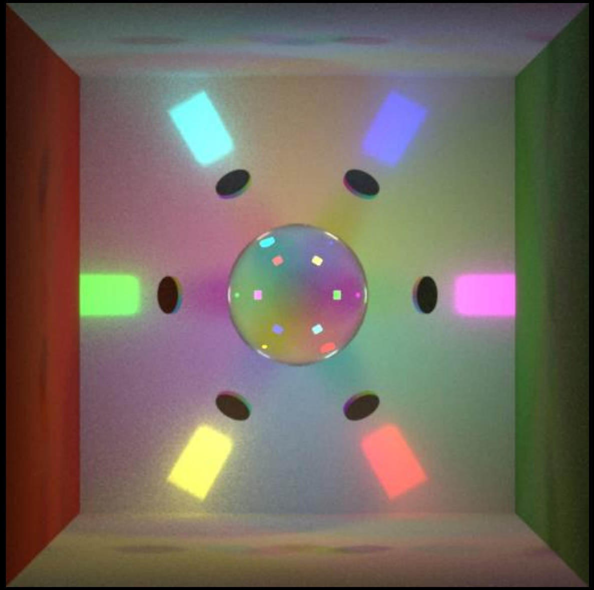

Target

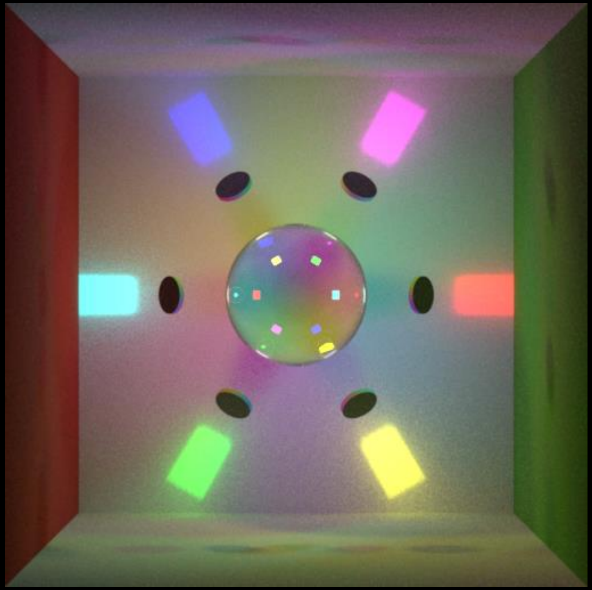

PRB[Vicini 2021]

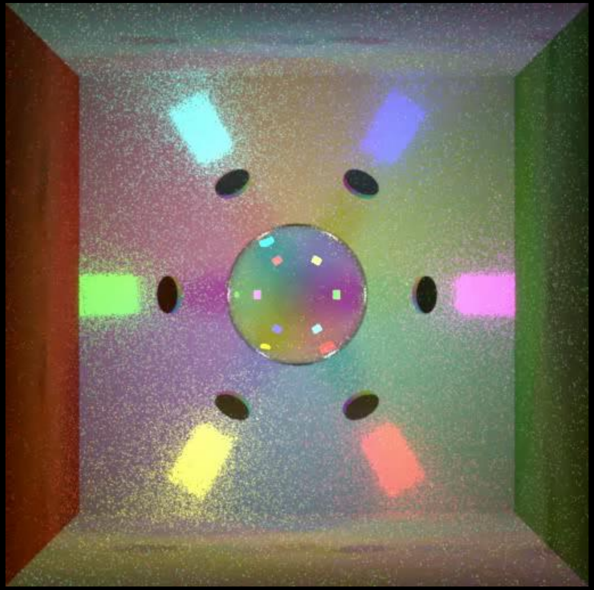

PRDPT[Fischer 2022]

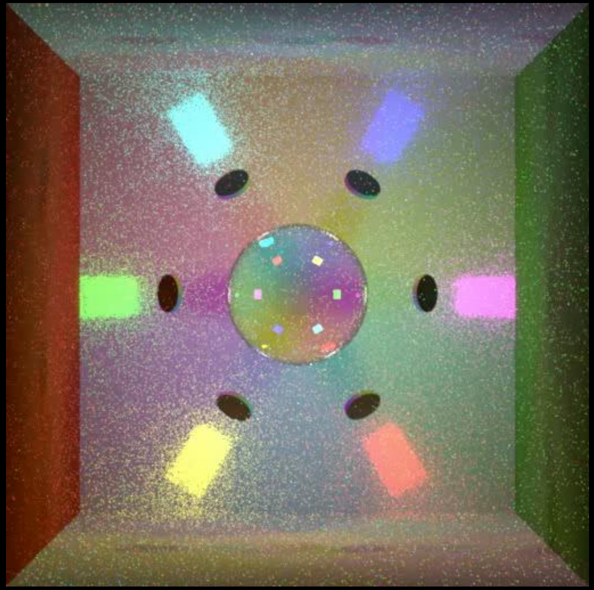

Ours

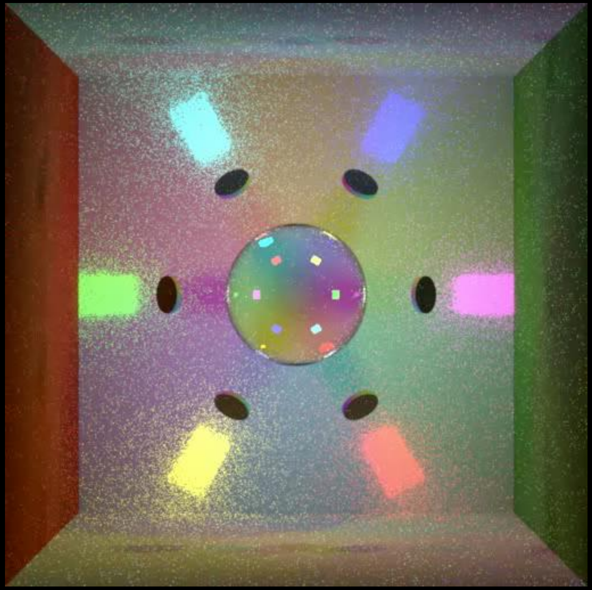

Shadow

Path Type: Shadow Path

Optimize parameter:

400 Spheres Occluder Translation

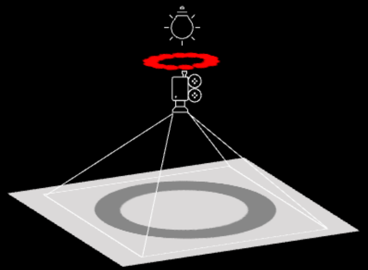

Init

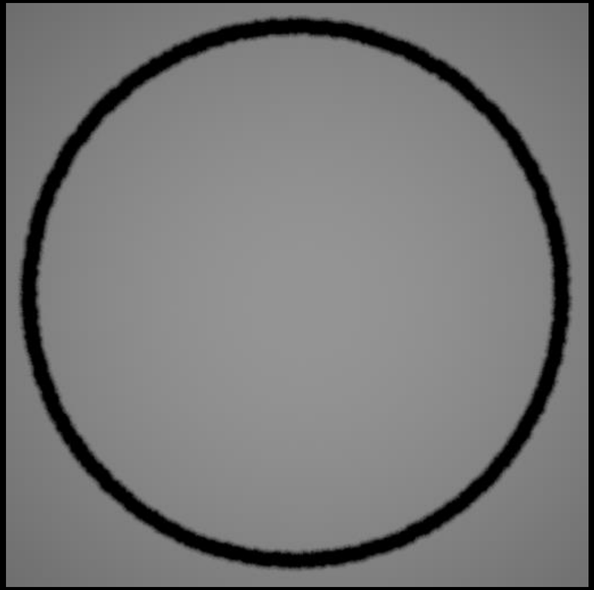

Target

PRB[Vicini 2021]

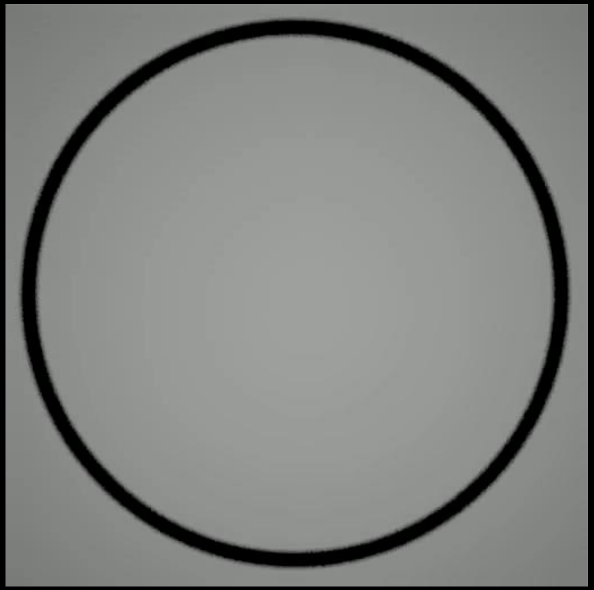

PRDPT[Fischer 2022]

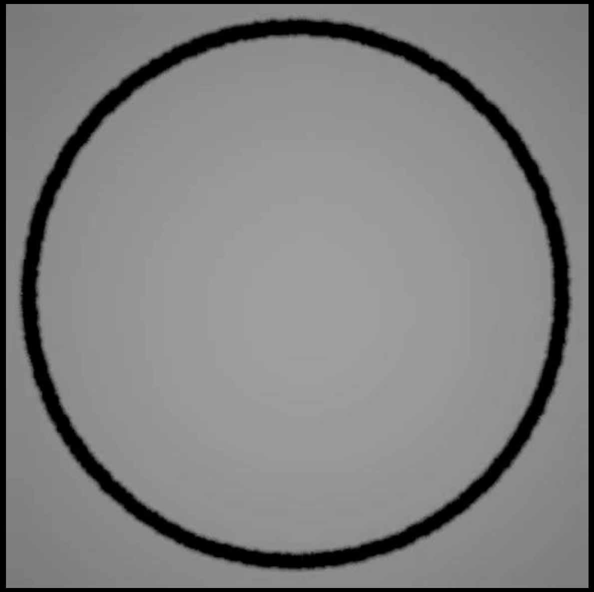

Ours

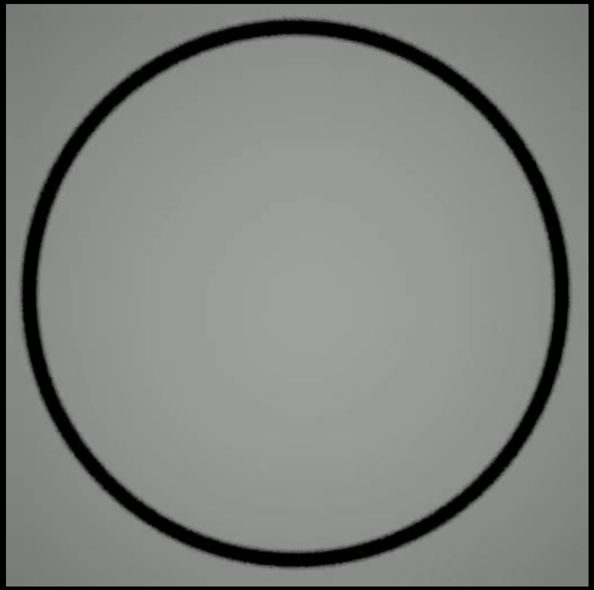

Results

For simple path

Gradient Visualization

θ: X-axis translation

Target Image

Initial

Traditional Gradients

Local and Sparse

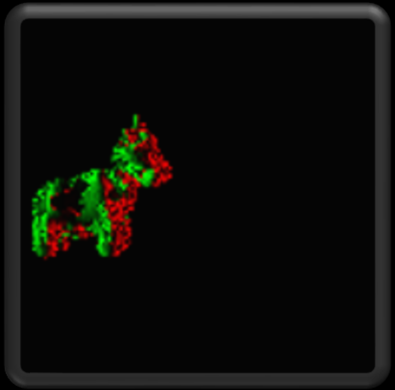

Nvdiffrast
Gradient

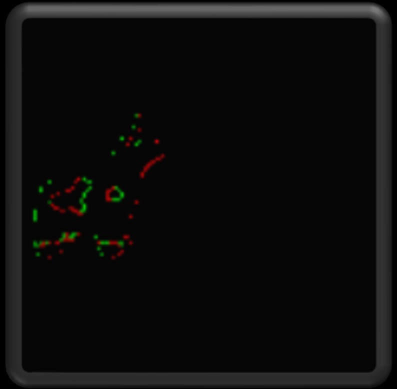

Nvdiffrast
Gradient
w/o shading

- Green: Indicates go right (right)
- Red: Indicates go left (wrong)

Our Gradients

Global and Dense

Human Pose Fitting

Initial Pose

Target Pose
(random)

PyTorch3D

Ours

Human Model: Skinned Multi-Person Linear Model (SMPL)
Optimize Parameter: 24 joints 3D rotation angle

Human Pose Fitting

Initial State

Target State
(random)

Nvdiffrast

Ours

Human Model: Skinned Multi-Person Linear Model (SMPL)
Optimize Parameter: 24 joints 3D rotation angle

Facial Expression Reconstruction

Initial State

Target State
(random)

PyTorch3D

Ours

Face Model: BlendShape with 11 base facial expression

Optimized Parameter: 11 weight parameters for each base expression

Facial Expression Reconstruction

Initial State

Target State
(random)

Nvdiffrast

Ours

Face Model: BlendShape with 11 base facial expression

Optimized Parameter: 11 weight parameters for each base expression

Furniture Layout

Initial State

Target State

PyTorch3D

Ours

Furniture Model: 3D-Front

Optimized Parameter: Furniture XZ-axis translation and Y-axis rotation

Furniture Layout

Initial State

Target State

Nvdiffrast

Ours

Furniture Model: 3D-Front

Optimized Parameter: Furniture XZ-axis translation and Y-axis rotation

Summary

- Contribution:
  - A novel pipeline for physically based differentiable rendering
  - A globally optimal transport based loss function
  - Formulation of extended path space manifolds to compute path derivatives
- Limitation & Future work:
  - Adapt EPSM to more types of path and scene representation
  - Derive both path geometry and color derivatives in EPSM formulation
  - Improve matching quality and efficiency

Thanks for Listening!