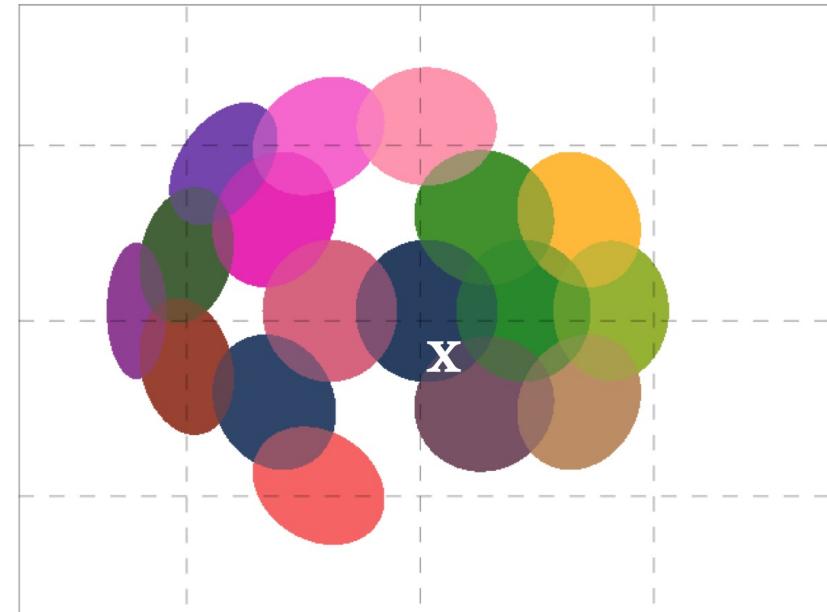


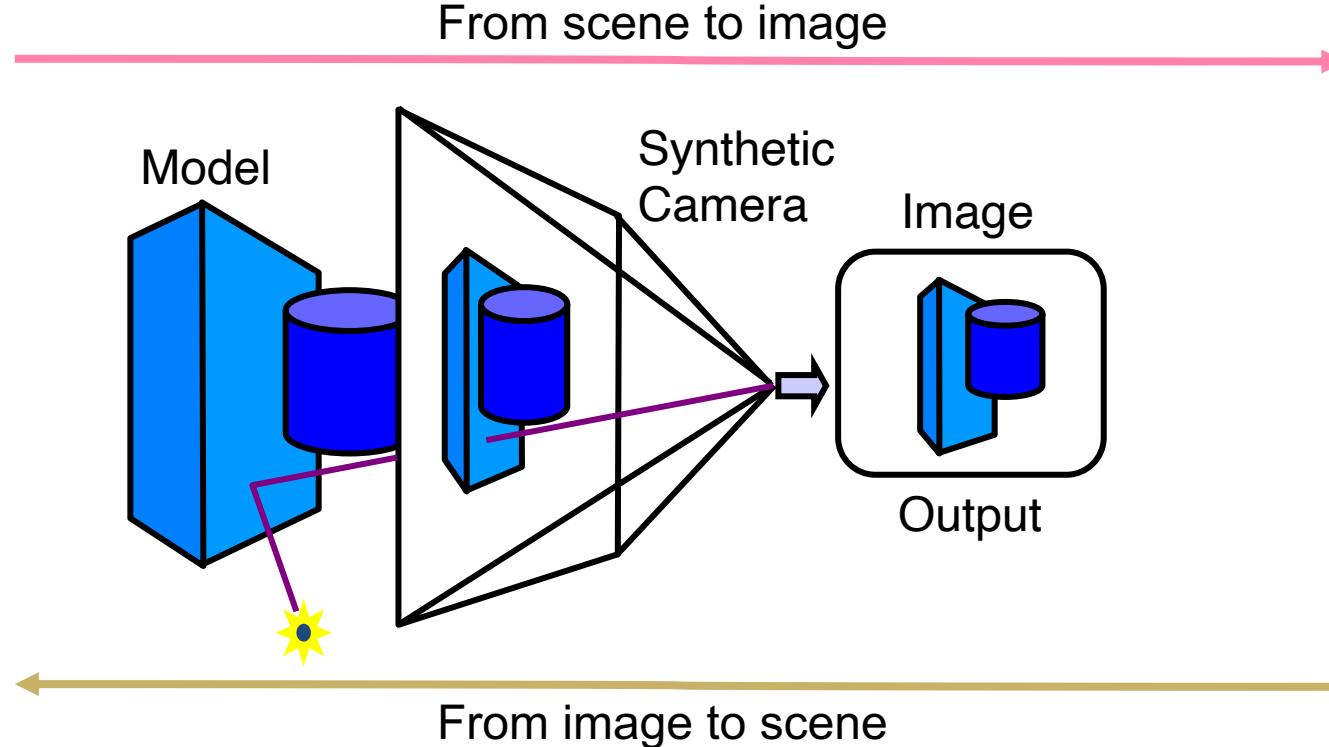
Inverse Rendering

Prof Cengiz Öztireli



Forward/ inverse rendering

Graphics is traditionally about forward rendering. That is what we have seen so far: scenes and their renderings. A recent problem in graphics is the inverse of this, i.e. going back from images to scenes via inverse rendering.



Forward / inverse rendering

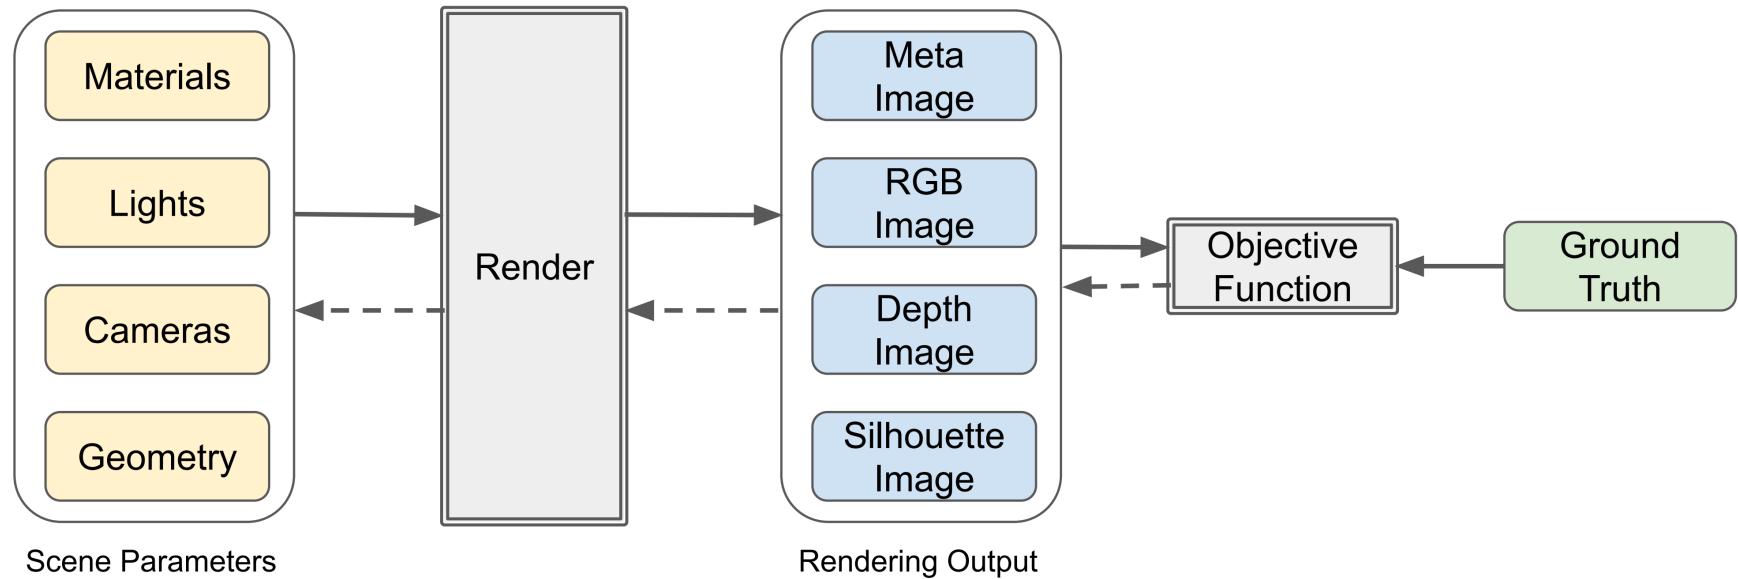
Models for
Geometry
Motion
Appearance
Lighting
Camera
...



Forward/ inverse rendering

Forward rendering means taking a scene definition with geometry, materials, cameras, and lights, and generating an image.

Inverse rendering starts from the rendered image, compares it to some form of ground truth image, and updates the parameters of the scene elements with the loss that compares the images.

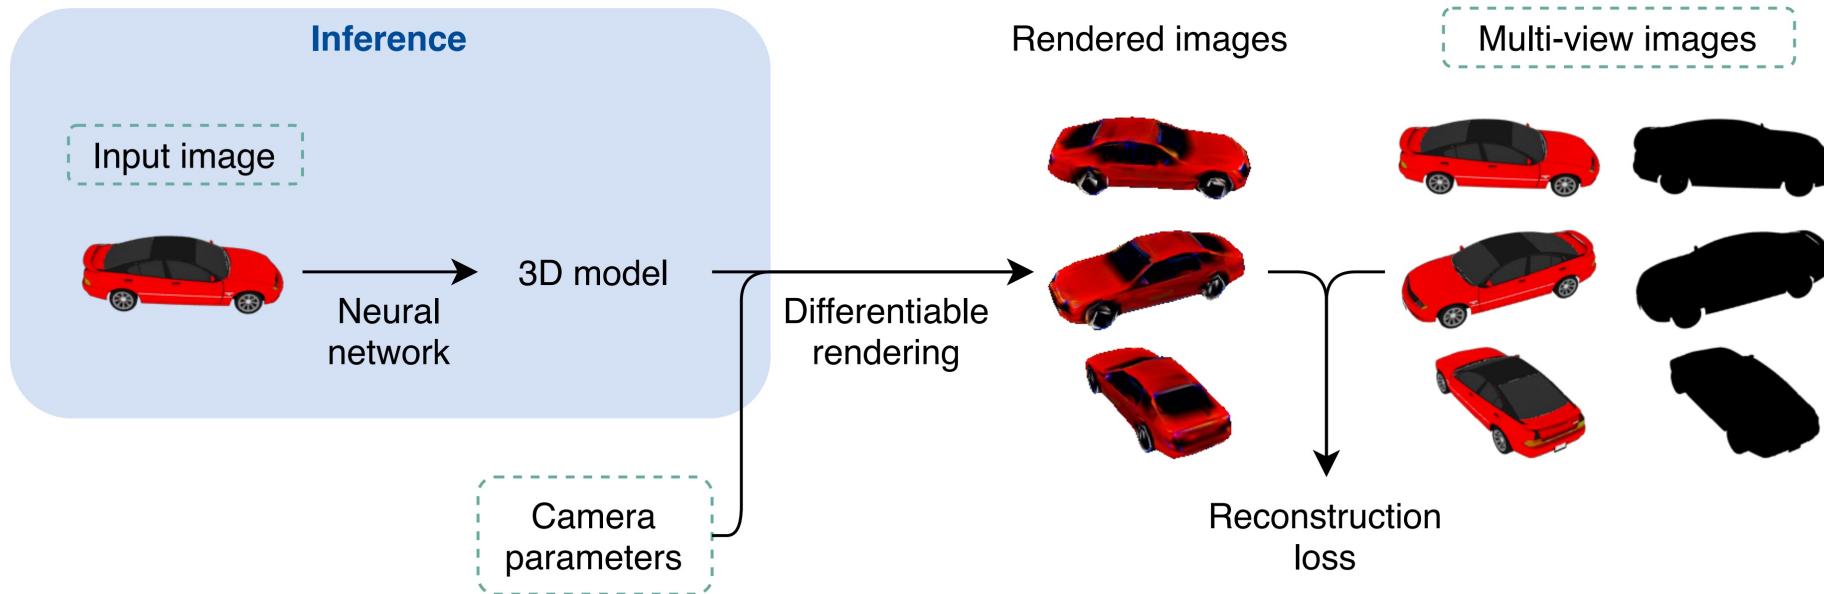


Forward/ inverse rendering

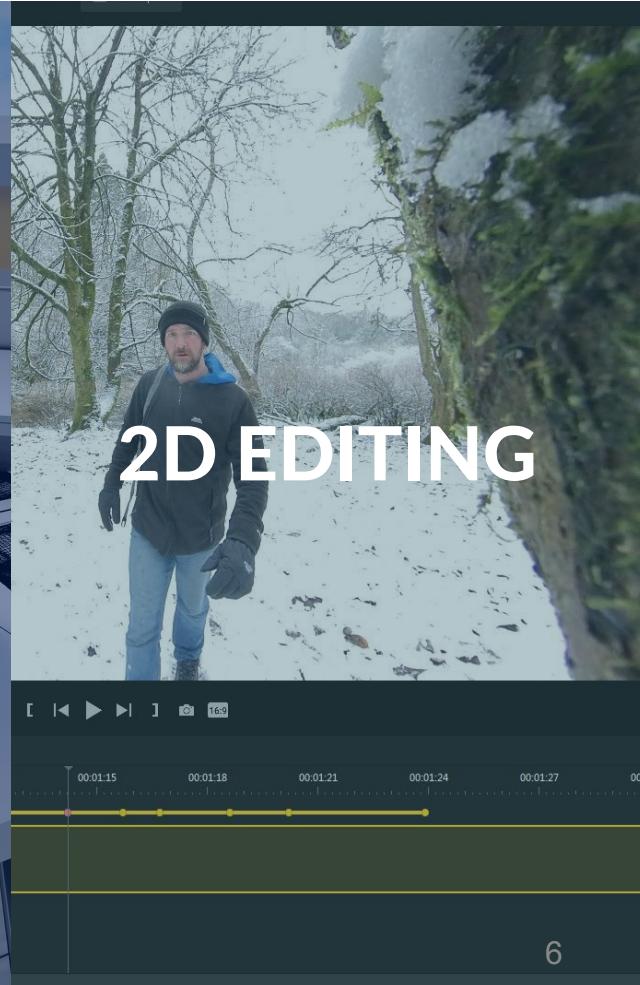
An application is geometry reconstruction from a single image.

At testing (inference) time, we can input an image into a neural network and get the corresponding 3D model. This network is optimized via rendering the 3D model generated from an image and comparing the rendered images with the ground truth images.

Example: 3D geometry reconstruction



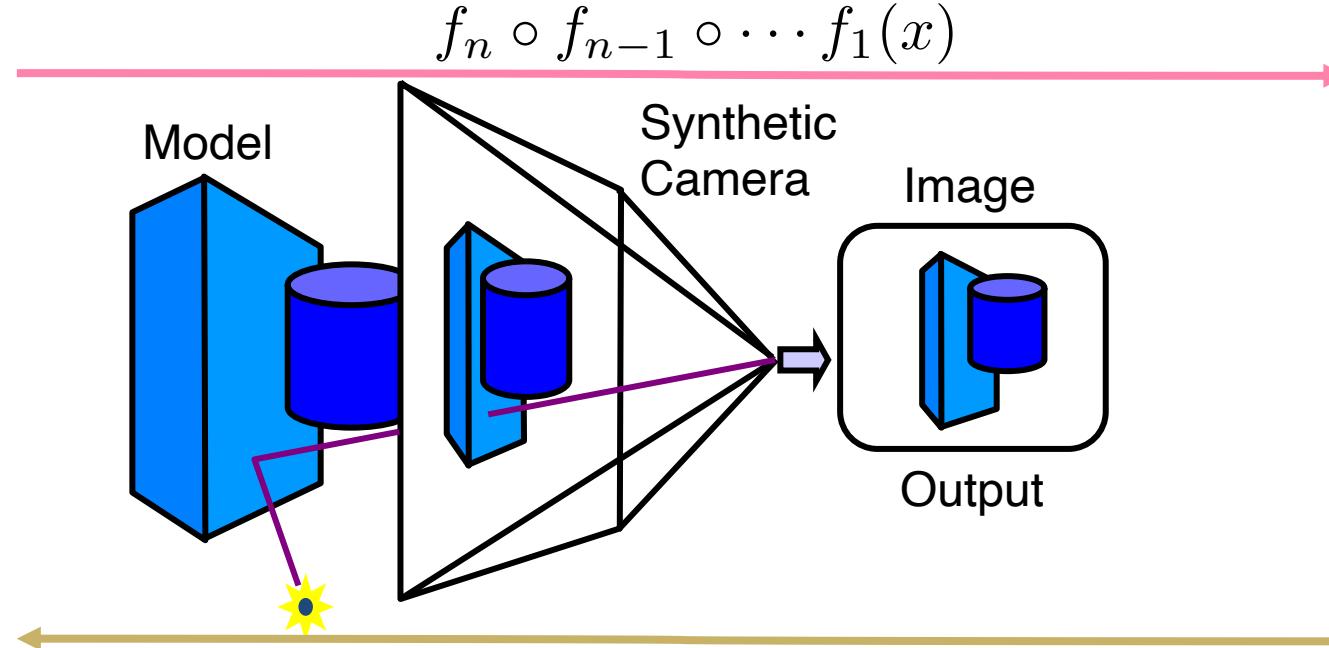
Inverse rendering has many applications.



Inverse rendering

Inverse rendering requires a non-linear optimization: the function from the scene parameters to rendered images is highly non-linear. Deep learning frameworks help us with this optimization, e.g. via stochastic gradient descent. There is one essential requirement for this to work: *we need differentiable rendering*.

- Deep learning to help: make everything differentiable



$$\frac{\partial f_n}{\partial f_{n-1}} \frac{\partial f_{n-1}}{\partial f_{n-2}} \cdots \frac{\partial f_1}{\partial x}$$

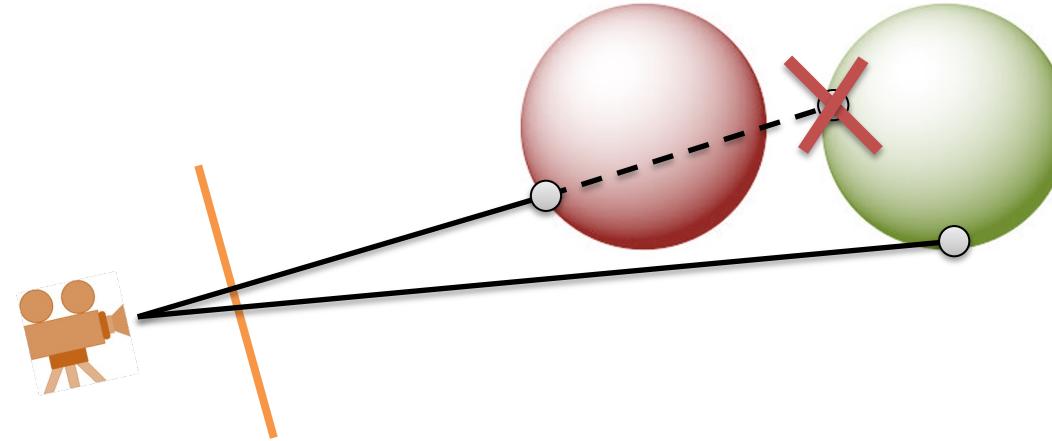
Inverse rendering

If you think about the rendering equation, everything is differentiable except one term: visibility.

Visibility is a binary function that can jump abruptly from 0 to 1 and 1 to 0.

This happens when e.g. an object occludes another one.

- Problem: Visibility is not differentiable



Visibility Term

Since visibility is discontinuous, the function we integrate is discontinuous.

Hence, we cannot exchange the order of derivatives with the integral.

This is, however, what we need for inverse rendering via optimizing scene parameters with gradient descent.

- **Problem with discontinuities and differentiability**

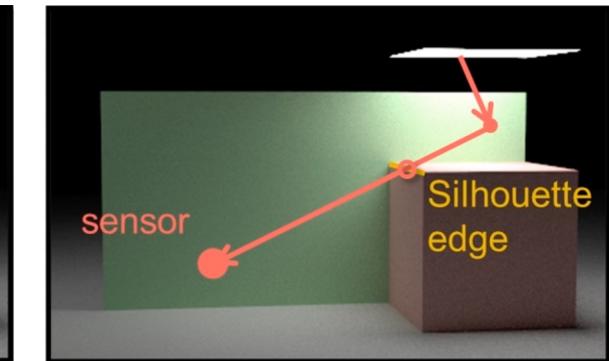
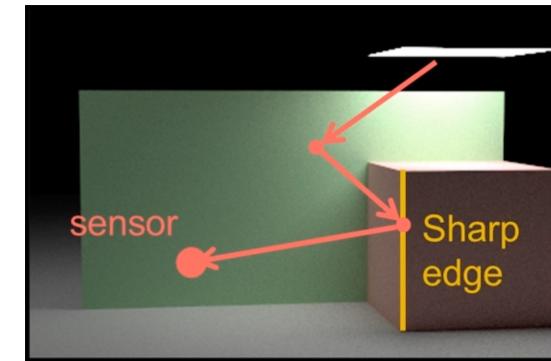
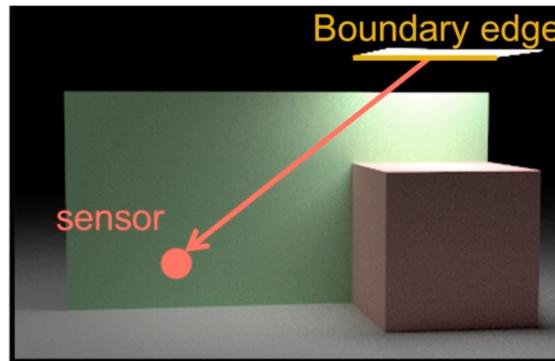
$$\frac{\partial}{\partial \pi} \int_A f(\mathbf{x}) d\mathbf{x} \neq \int_A \frac{\partial}{\partial \pi} f(\mathbf{x}) d\mathbf{x}$$

Visibility Term

Example discontinuities of visibility occurs for:

- at the boundary of a light source, where light suddenly stops reaching the sensor point,
- at a sharp edge of an object, where light suddenly stops bouncing,
- at a silhouette edge, where an object is occluding a light source when viewed from the sensor.

• Sources of discontinuities due to visibility



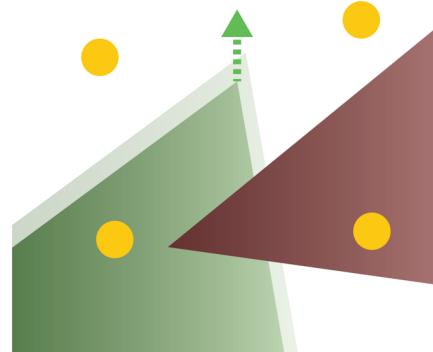
Differentiable Visibility

No need to know the details of differentiable visibility techniques.

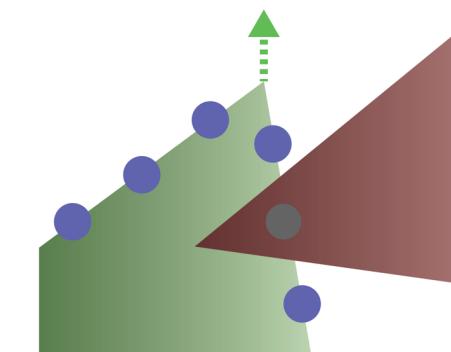
There are a few techniques to provide a continuous approximation of visibility.

One of them is densely sampling polygon edges in the meshes in a scene, as for meshes those are where discontinuities appear.

- Edge-sampling for accurate gradients



Standard sampling

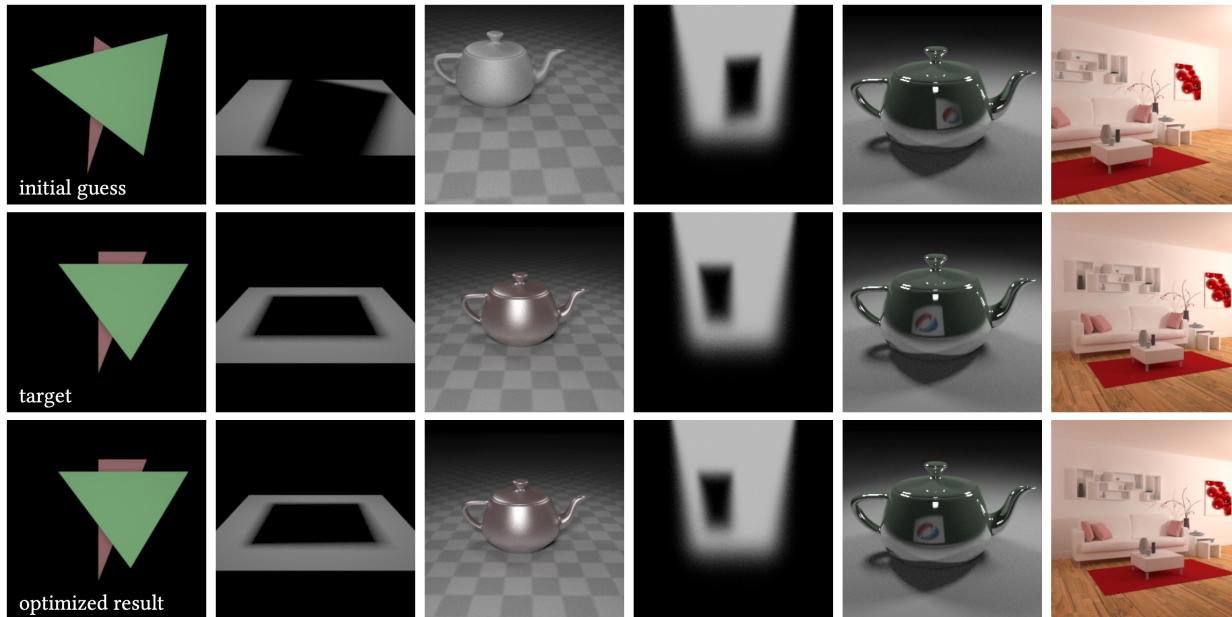


Edge sampling

Li et al., Differentiable Monte Carlo Ray Tracing through Edge Sampling, 2018.

Differentiable Visibility

- Edge-sampling for accurate gradients

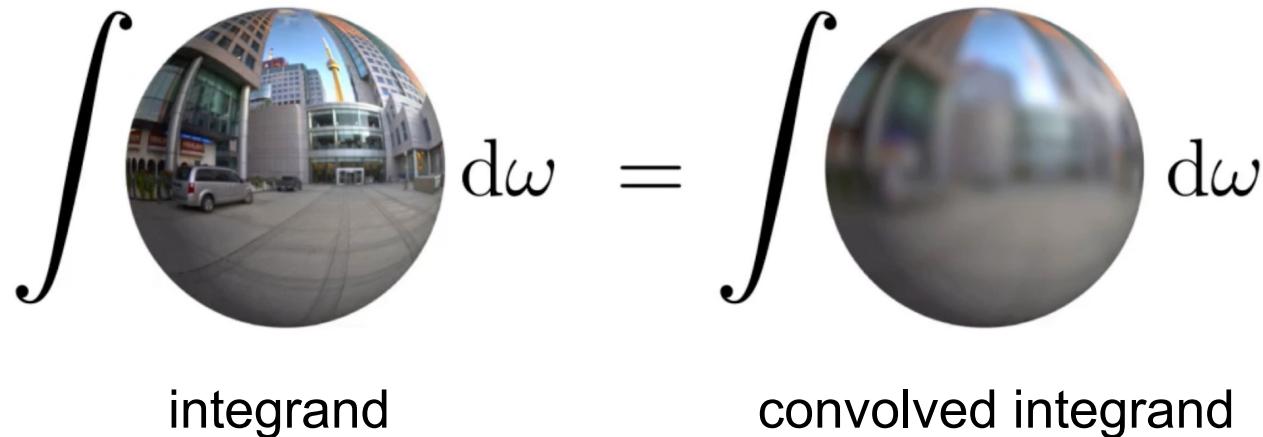


Li et al., Differentiable Monte Carlo Ray Tracing through Edge Sampling, 2018.

Differentiable Visibility

Another idea is convolving the function we are integrating to smooth it to get rid of discontinuities.

- Re-parametrizing discontinuities



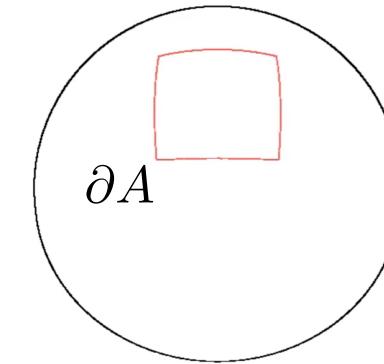
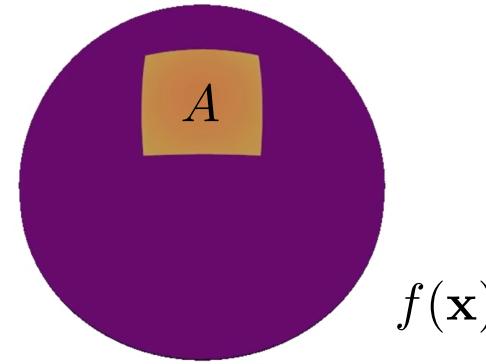
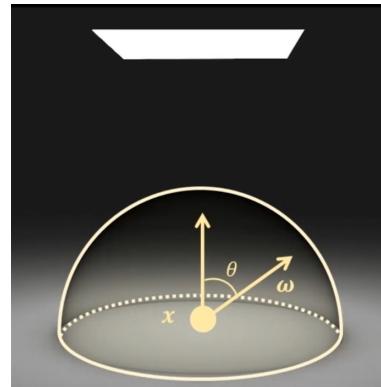
Loubet et al., Reparameterizing discontinuous integrands for differentiable rendering, 2019.

Differentiable Visibility

Yet another idea is utilising a general theorem that provides the integral of a discontinuous function. This theorem introduces an additional integral (last term below) over the boundary of the area we are originally integrating over.

- Reynolds transport theorem

$$\frac{\partial}{\partial \pi} \int_A f(\mathbf{x}) d\mathbf{x} = \int_A \frac{\partial}{\partial \pi} f(\mathbf{x}) d\mathbf{x} + \int_{\partial A} g(\mathbf{l}) d\mathbf{l}$$



Zhang et al., Path-Space Differentiable Rendering, 2020.

Frameworks

- Automatic differentiation
- Fast & parallel computation
- Active development

TensorFlow
Graphics



Feedback & Actions

We will have sessions for **going through examples** for the ones who could not get enough supervision.

- These replace office hours.
- If you cannot make it, do let me know.

No need to know more than what is covered in the lecture notes (annotated slides).

Make sure you study with the **lecture notes = annotated slides with details**