Scene Understanding — Differentiable Graphics

Dr Fangcheng Zhong

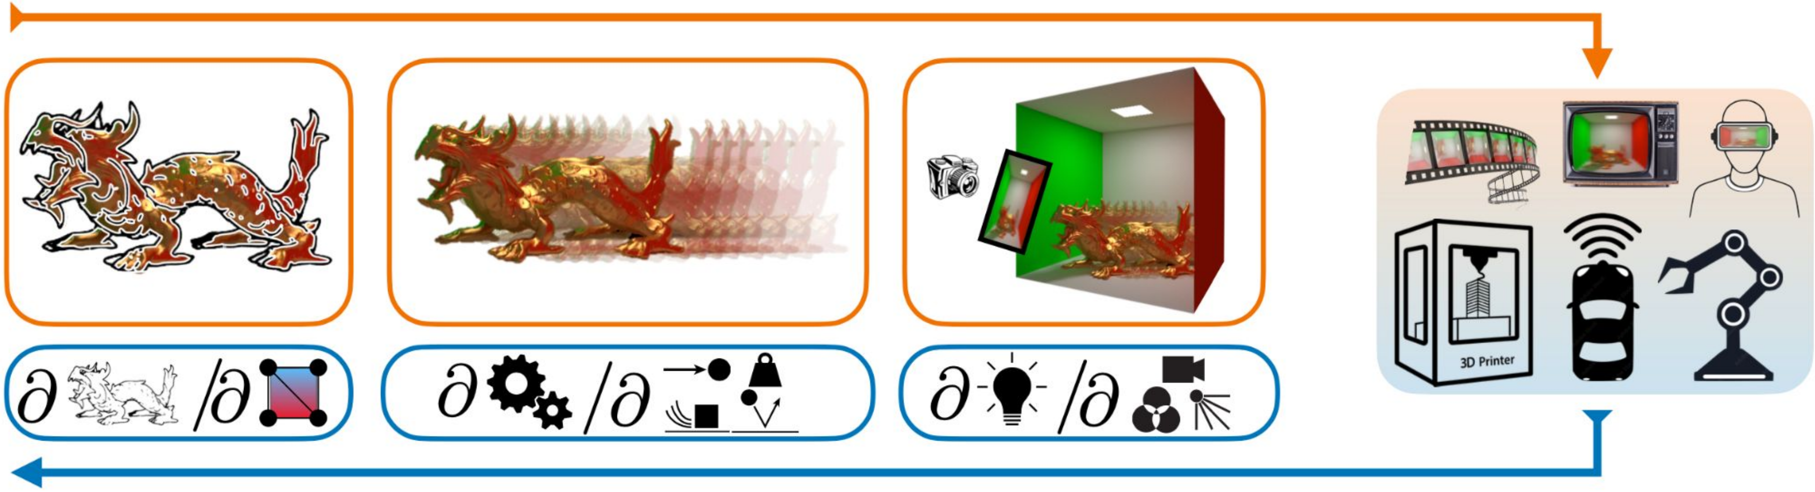

UNIVERSITY OF CAMBRIDGE

Outline

- Inverse graphics
- Differentiable surface rendering
  - Differentiable rasterization
  - Physically-based differentiable rendering
- Non-surface representations
  - Light fields
  - Neural radiance fields
  - Differentiable volume rendering

Inverse Graphics

Rendering

Inverse rendering

Inverse Graphics

- Traditional approaches
  - SLAM, SfM
  - Light probes, structured light
- Data-driven approaches

Differentiable Rendering

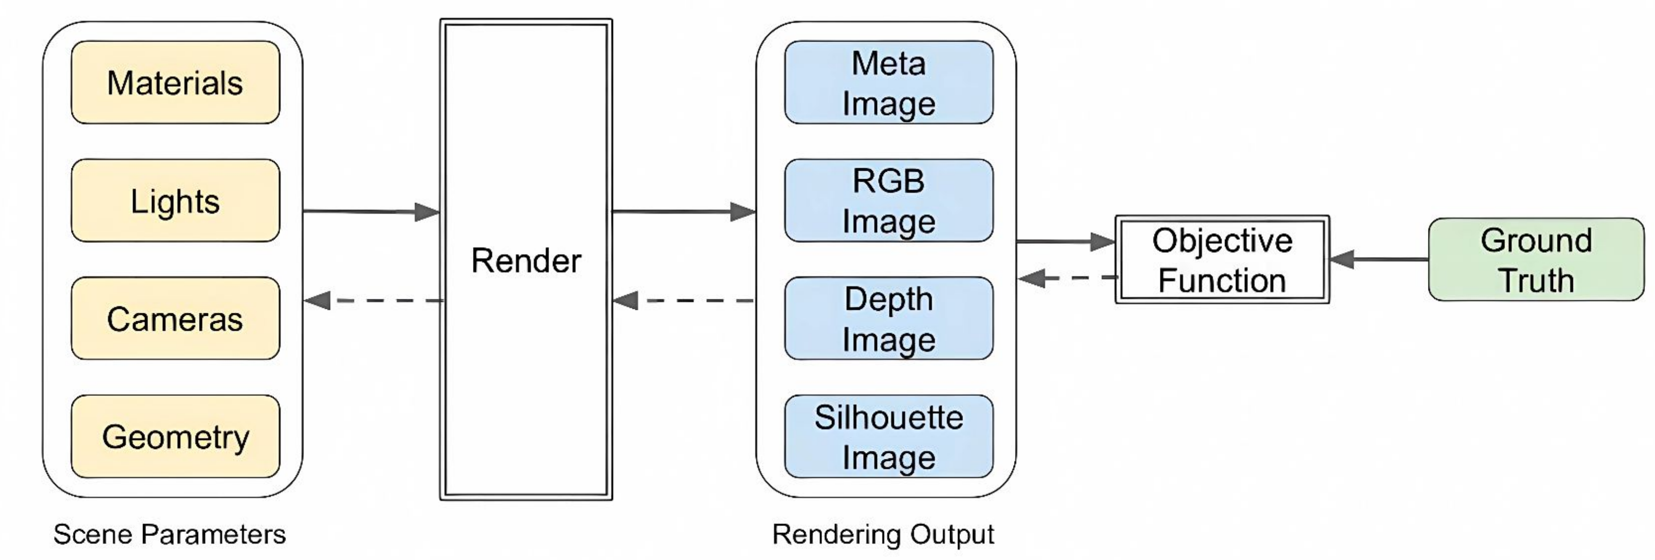

generalised reprojection error minimisation

Differentiable Rendering

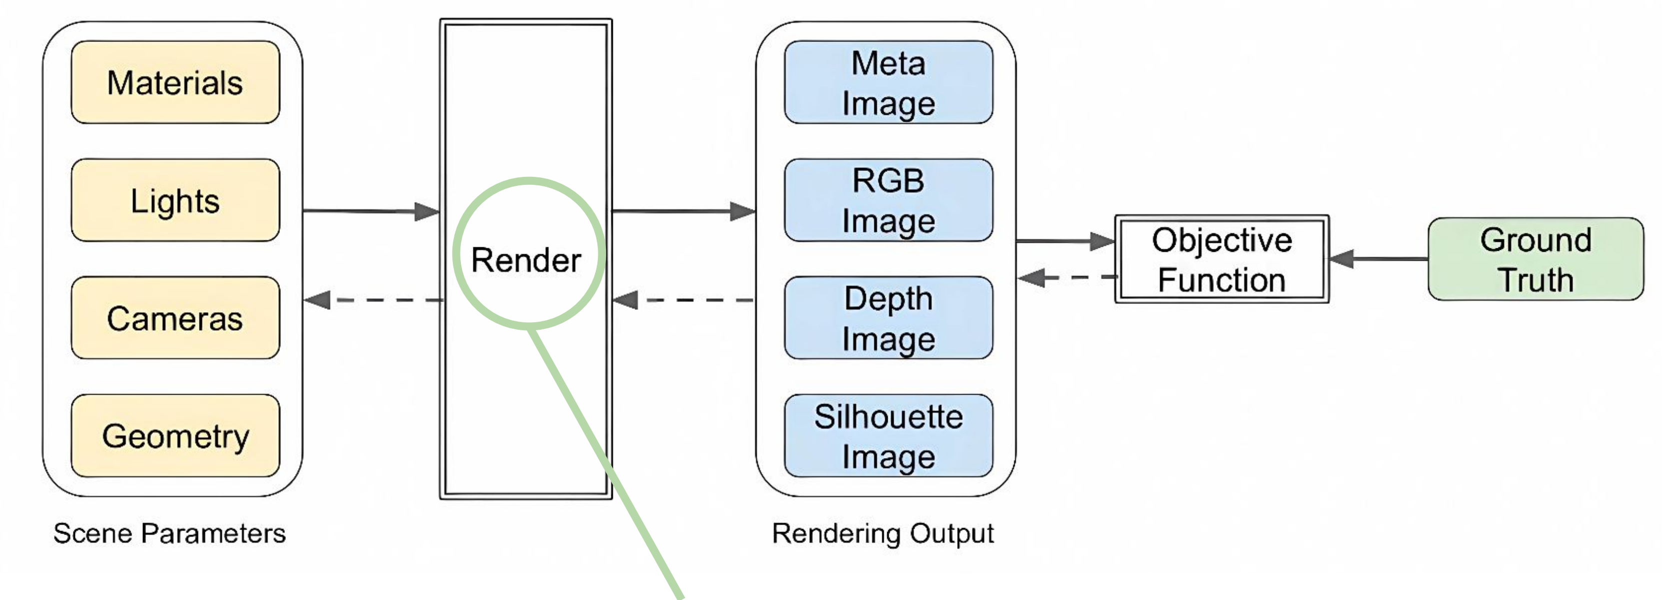

Derive useful gradients in rendering

Differentiable Graphics

Everything differentiable can be integrated!

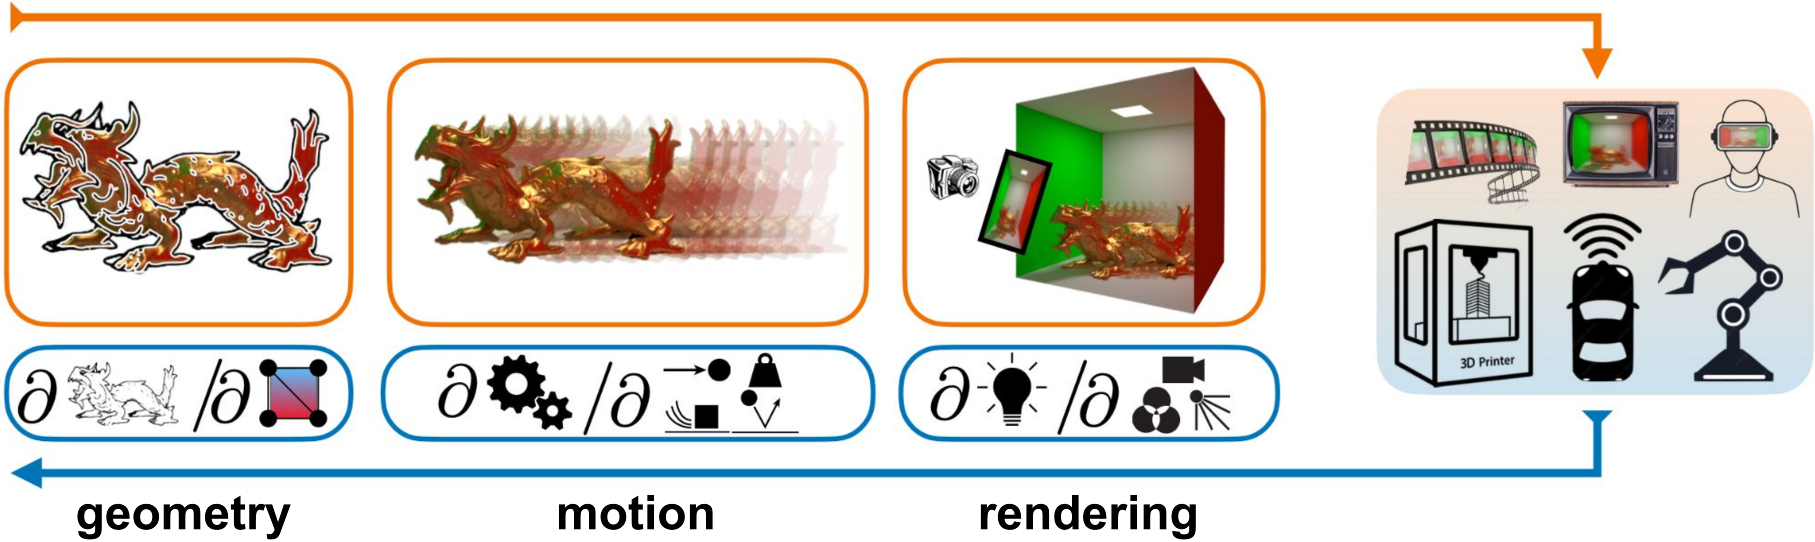

Differentiable Graphics

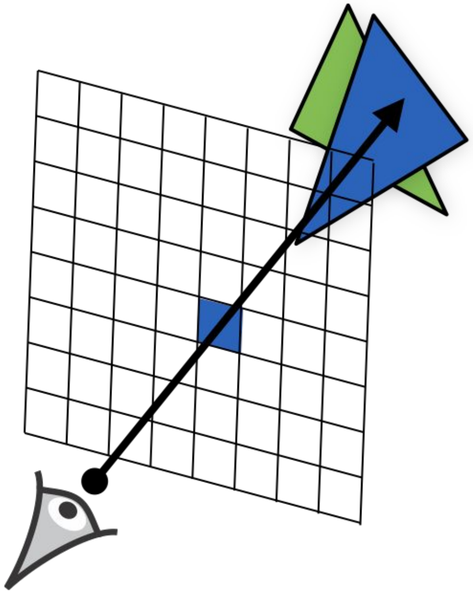

Visibility has no well-defined gradient

OpenGL Rendering Pipeline

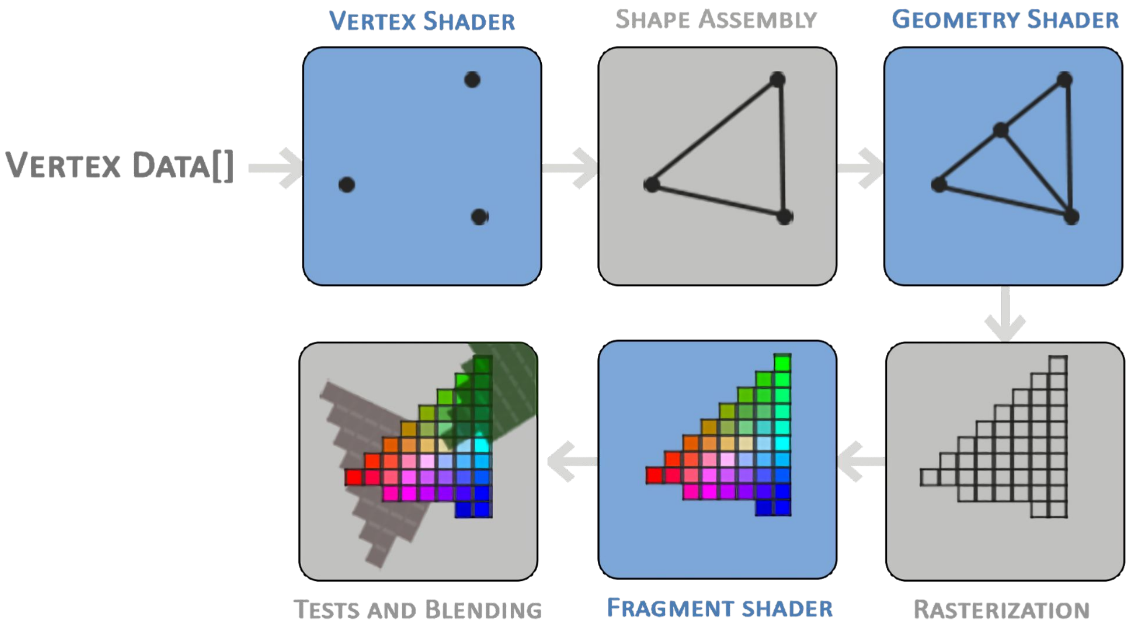

Differentiable Rasterization

Rasterization has no gradient

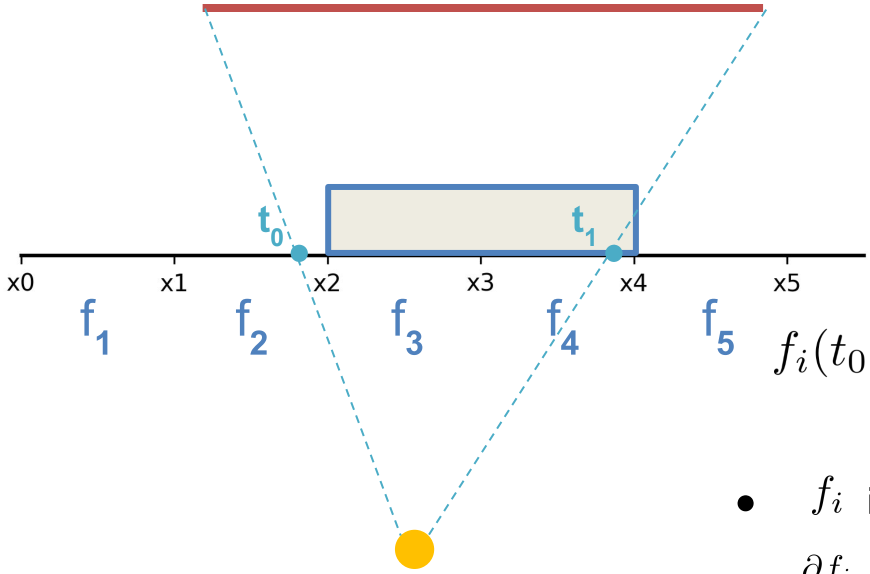

$$f_i(t_0, t_1) = \begin{cases} 1 & \text{if } [t_0, t_1] \cap [x_i, x_{i-1}] < \frac{x_i - x_{i-1}}{2} \\ 0 & \text{otherwise} \end{cases}$$

- f_i is the visibility of [red line] at the i-th pixel location
- $\frac{\partial f_i}{\partial t_j}$ is zero or undefined everywhere

UNIVERSITY OF CAMBRIDGE

Differentiable Rasterization

Z-buffering has no gradient

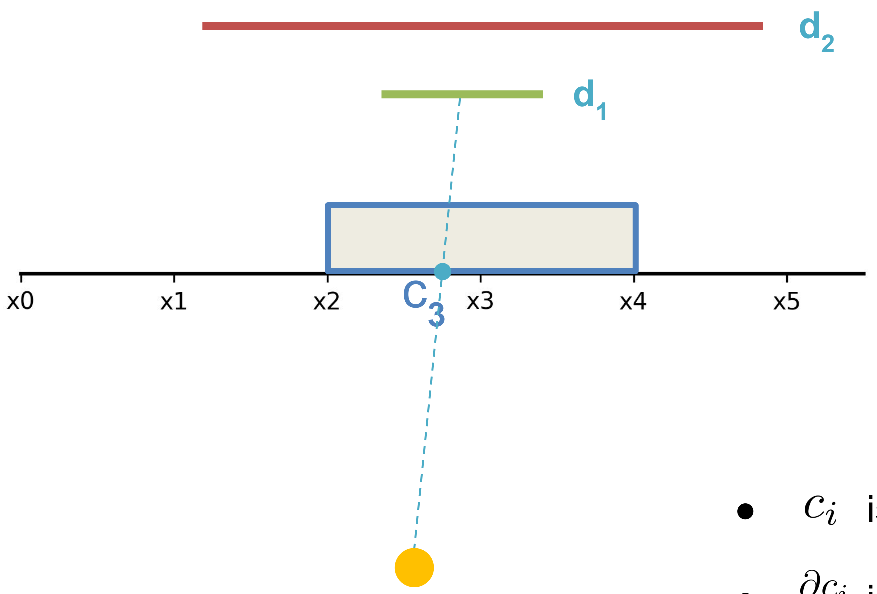

$$c_i(d_1, d_2) = \begin{cases} \text{green} & \text{if } d_1 < d_2 \\ \text{red} & \text{otherwise} \end{cases}$$

- c_i is the colour of i-th pixel location (if visible)
- $\frac{\partial c_i}{\partial d_j}$ is zero or undefined everywhere

Differentiable Rasterization

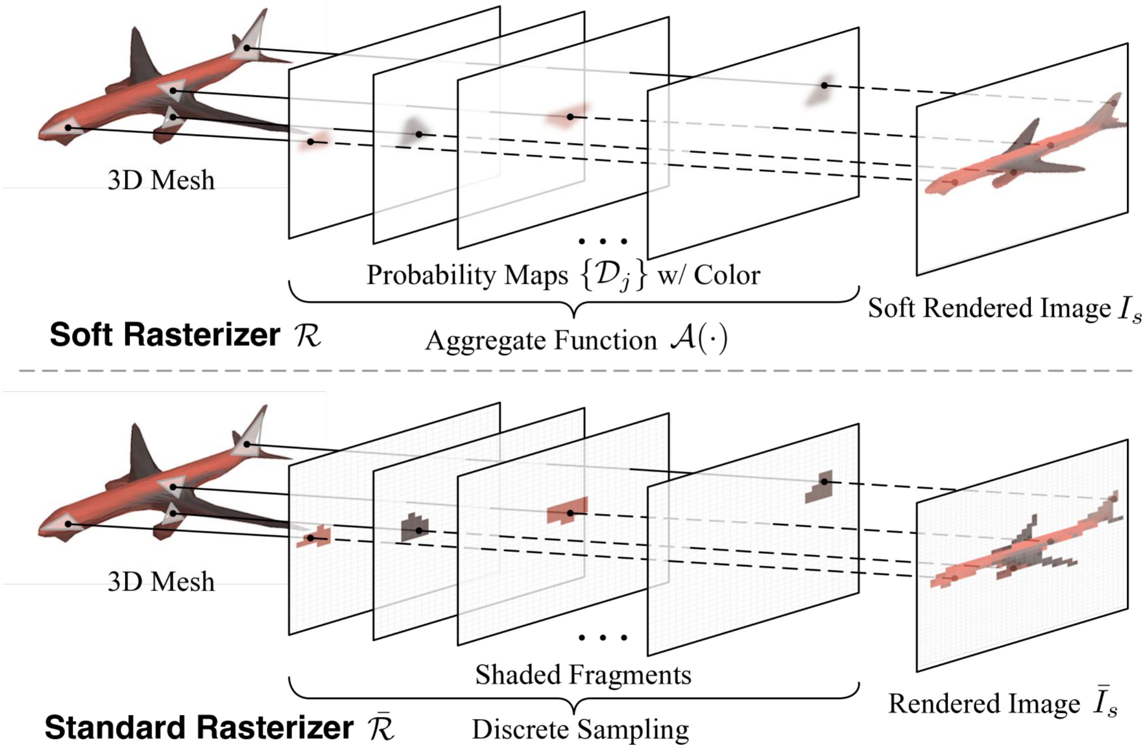

UNIVERSITY OF CAMBRIDGE

Liu, Shichen, et al. "Soft rasterizer: A differentiable renderer for image-based 3d reasoning." *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019.

Differentiable Rasterization

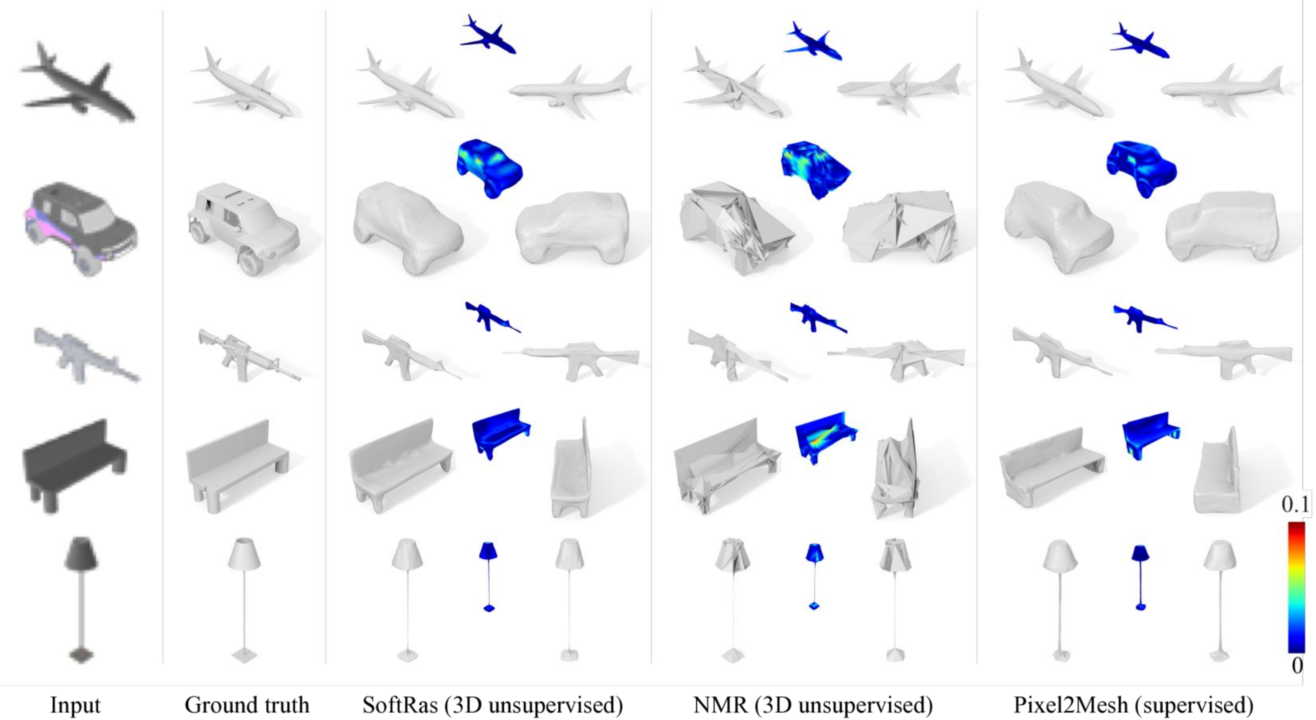

Liu, Shichen, et al. "Soft rasterizer: A differentiable renderer for image-based 3d reasoning." *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019.

Differentiable Rasterization

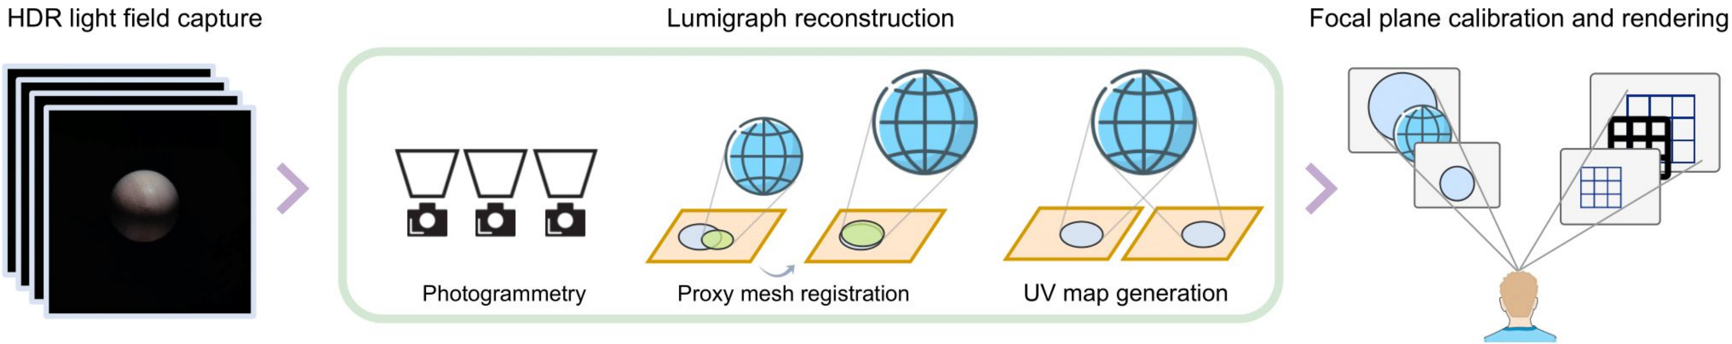

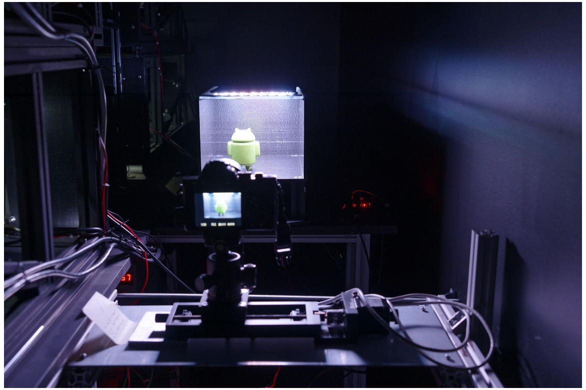

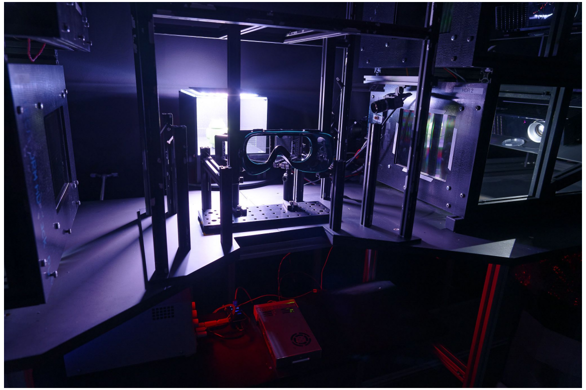

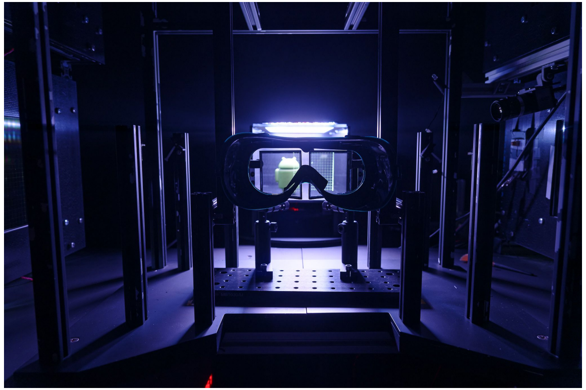

Fangcheng Zhong, Akshay Jindal, Ali Özgür Yöntem, Param Hanji, Simon J. Watt, and Rafał K. Mantiuk. 2021. Reproducing Reality with a High-Dynamic-Range Multi-Focal Stereo Display. ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia, Journal Track), 2021

Physically-based Rendering

The rendering equation

$$L_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_\Omega f_r(\mathbf{x}, \omega_i, \omega_o) L_i(\mathbf{x}, \omega_i) \cos(\theta) d\omega_i$$

discontinuous!

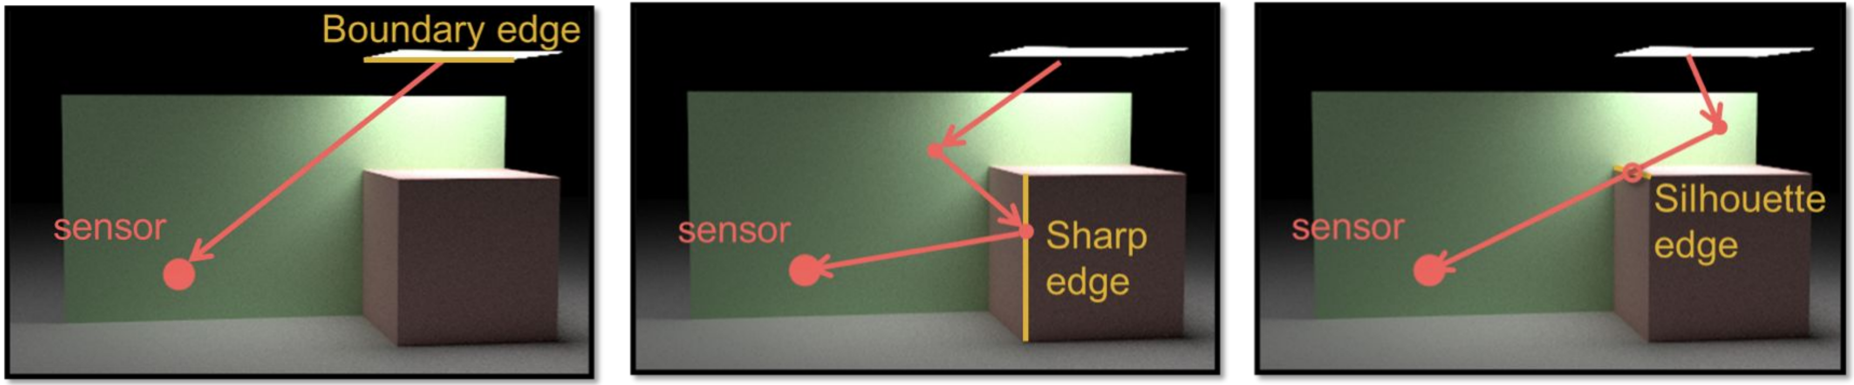

UNIVERSITY OF CAMBRIDGE

Physically-based Differentiable Rendering

Differentiating the rendering equation

$$L_o(\mathbf{x}, \omega_o; \Theta) = L_e(\mathbf{x}, \omega_o; \Theta) + \int_\Omega f_r(\mathbf{x}, \omega_i, \omega_o; \Theta) L_i(\mathbf{x}, \omega_i; \Theta) \cos(\theta) d\omega_i$$

scene parameters

$$\frac{d}{d\Theta} L_o(\mathbf{x}, \omega_o; \Theta) \neq \frac{d}{d\Theta} L_e(\mathbf{x}, \omega_o; \Theta) + \int_\Omega \frac{d}{d\Theta} f_r(\mathbf{x}, \omega_i, \omega_o; \Theta) L_i(\mathbf{x}, \omega_i; \Theta) \cos(\theta) d\omega_i$$

Only true when the integrand is continuous

Physically-based Differentiable Rendering

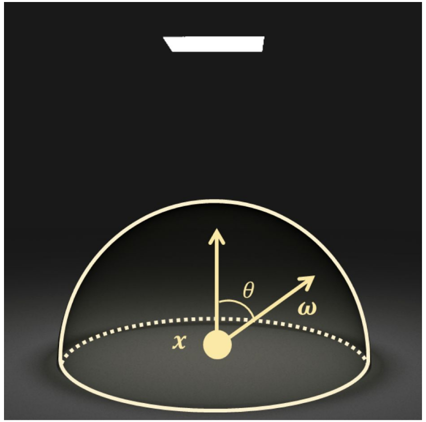

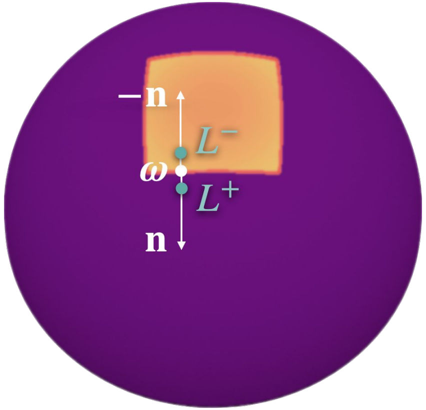

boundary of the integration domain

discontinuous boundary of L_i

$$\frac{d}{d\Theta}L_o = \frac{d}{d\Theta}L_e + \int_{\Omega}\frac{d}{d\Theta}f_r\,L_i\,\cos\,d\omega_i + \int_{\partial\Omega\cup\Omega^*}\mathbf{v}\cdot\mathbf{n}\,(L_i^- - L_i^+)\,f_r\,\cos\,dS$$

differential rendering equation

movement of S w.r.t. theta in the normal direction

UNIVERSITY OF CAMBRIDGE

Physically-based Differentiable Rendering

Leibniz integral rule

General form:

$$\frac{d}{dx}\int_{a(x)}^{b(x)} f(x,t)\,dt = f(x,b(x))\frac{db}{dx} - f(x,a(x))\frac{da}{dx} + \int_{a(x)}^{b(x)} \frac{\partial}{\partial x} f(x,t)\,dt$$

If $i)$ $f(x,t)$ and its partial derivative $f_x(x,t)$ are continuous; $ii)$ a and b are constant independent of x,

$$\frac{d}{dx}\int_{a}^{b} f(x,t)\,dt = \int_{a}^{b} \frac{\partial}{\partial x} f(x,t)\,dt$$

Physically-based Differentiable Rendering

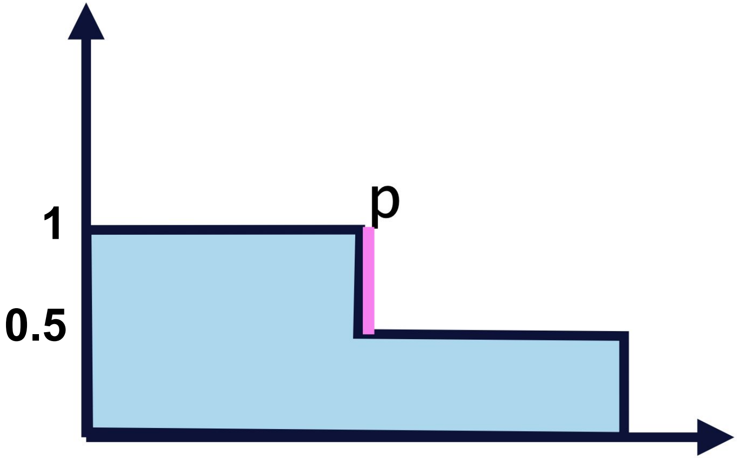

$$\frac{d}{dp}\int_a^b t < p\,?\,1 : 0.5\,dt$$

What if the integrand has sharp discontinuities (e.g. visibility) ?

Physically-based Differentiable Rendering

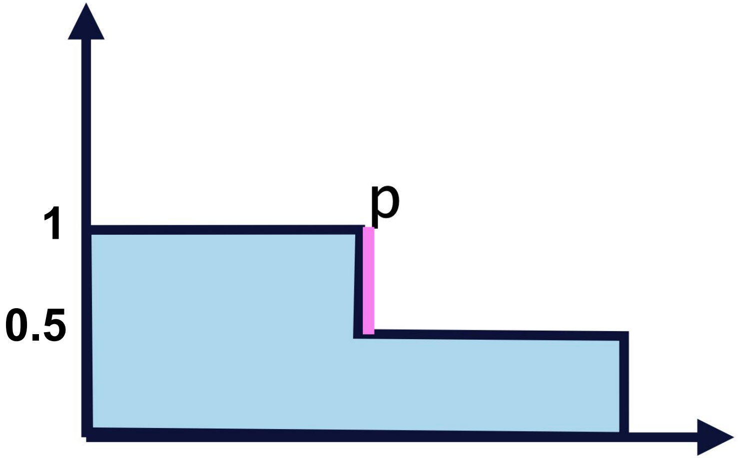

$$\frac{d}{dp}\int_a^b t < p \,?\, 1 : 0.5 \, dt$$

$$= \frac{d}{dp}\int_a^p 1 \, dt + \frac{d}{dp}\int_p^b 0.5 \, dt$$

$$= 1 - 0.5$$

$$= 0.5$$

Can apply the Leibniz rule to each term

Physically-based Differentiable Rendering

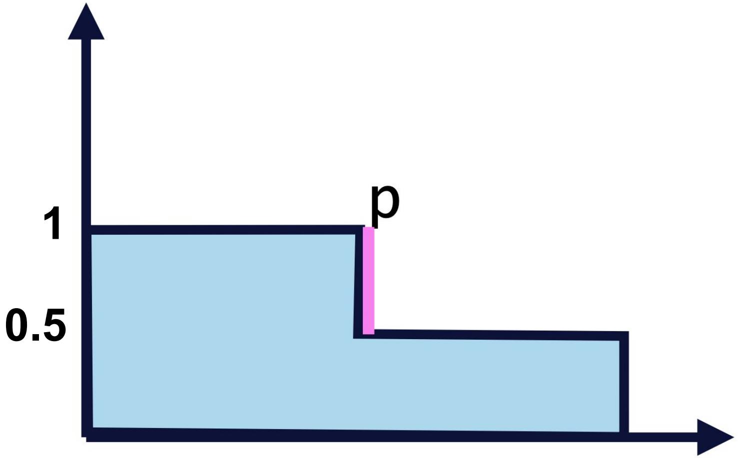

$$\frac{d}{dp}\int_a^b t < p\,?\,1 : 0.5\,dt$$

$$= \frac{d}{dp}\int_a^p 1\,dt + \frac{d}{dp}\int_p^b 0.5\,dt$$

$$= 1 - 0.5$$

$$= 0.5$$

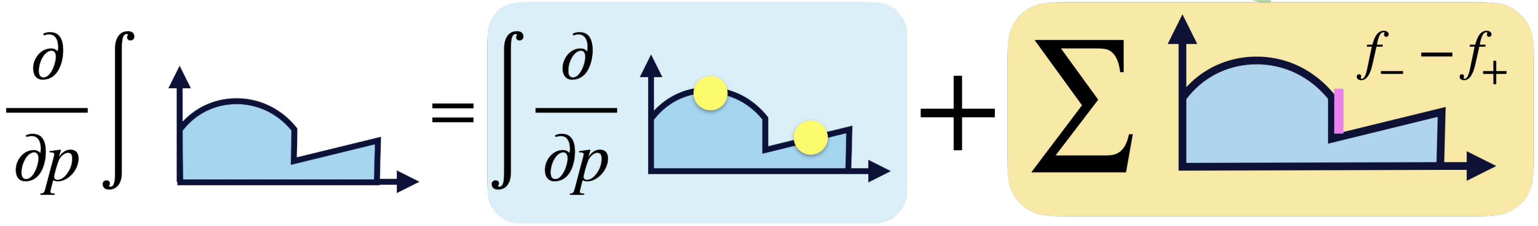

Physically-based Differentiable Rendering

Reynolds transport theorem

$$\frac{d}{dt}\int_{\Omega(t)} f(\mathbf{x},t)\,dV(\mathbf{x}) = \int_{\Omega(t)} \frac{\partial}{\partial t} f(\mathbf{x},t)\,dV + \int_{\partial\Omega(t)} f(\mathbf{x},t)\mathbf{v}\cdot\mathbf{n}\,dA$$

generalisation of the Leibniz integral rule to higher dimensions

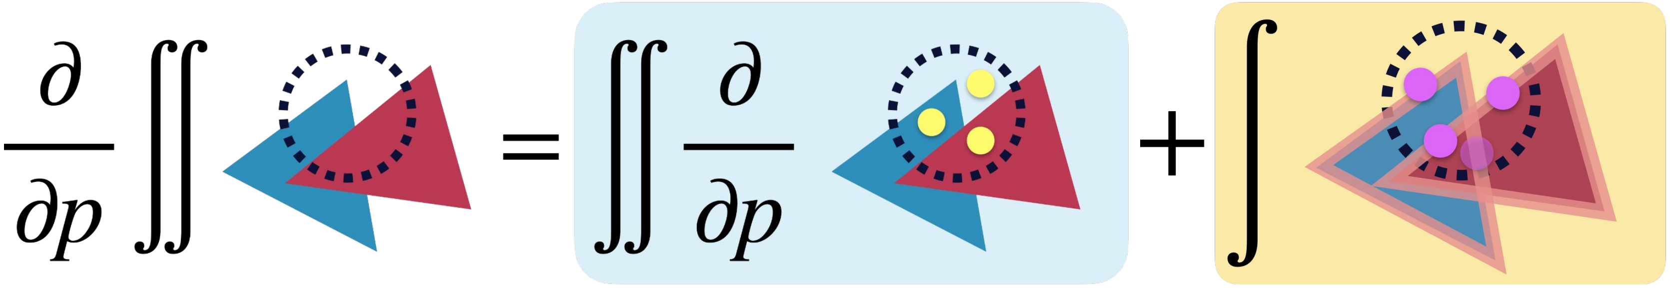

UNIVERSITY OF CAMBRIDGE

Physically-based Differentiable Rendering

Loubet, Guillaume, Nicolas Holzschuch, and Wenzel Jakob. "Reparameterizing discontinuous integrands for differentiable rendering." *ACM Transactions on Graphics (TOG)* 38.6 (2019): 1-14.

Light Fields

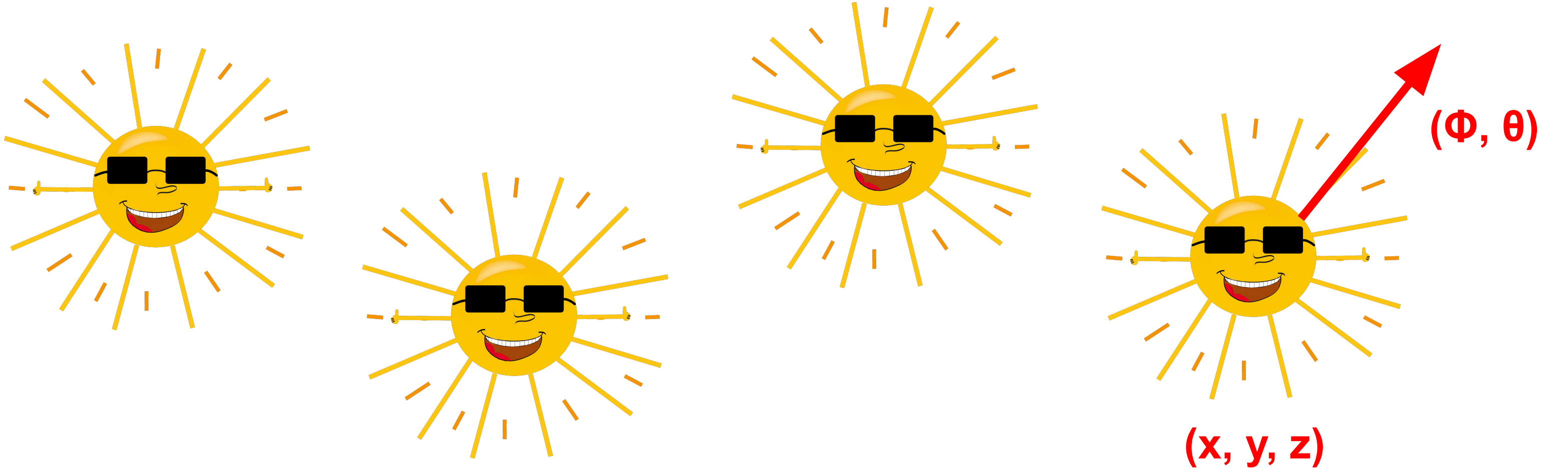

$$f(x, y, z, \phi, \theta)$$

Incident radiance at an arbitrary location from an arbitrary direction

Light Fields

Broxton, Michael, et al. "Immersive light field video with a layered mesh representation." *ACM Transactions on Graphics (TOG)* 39.4 (2020): 86-1. https://augmentedperception.github.io/deepviewvideo/

Neural Radiance Field (NeRF)

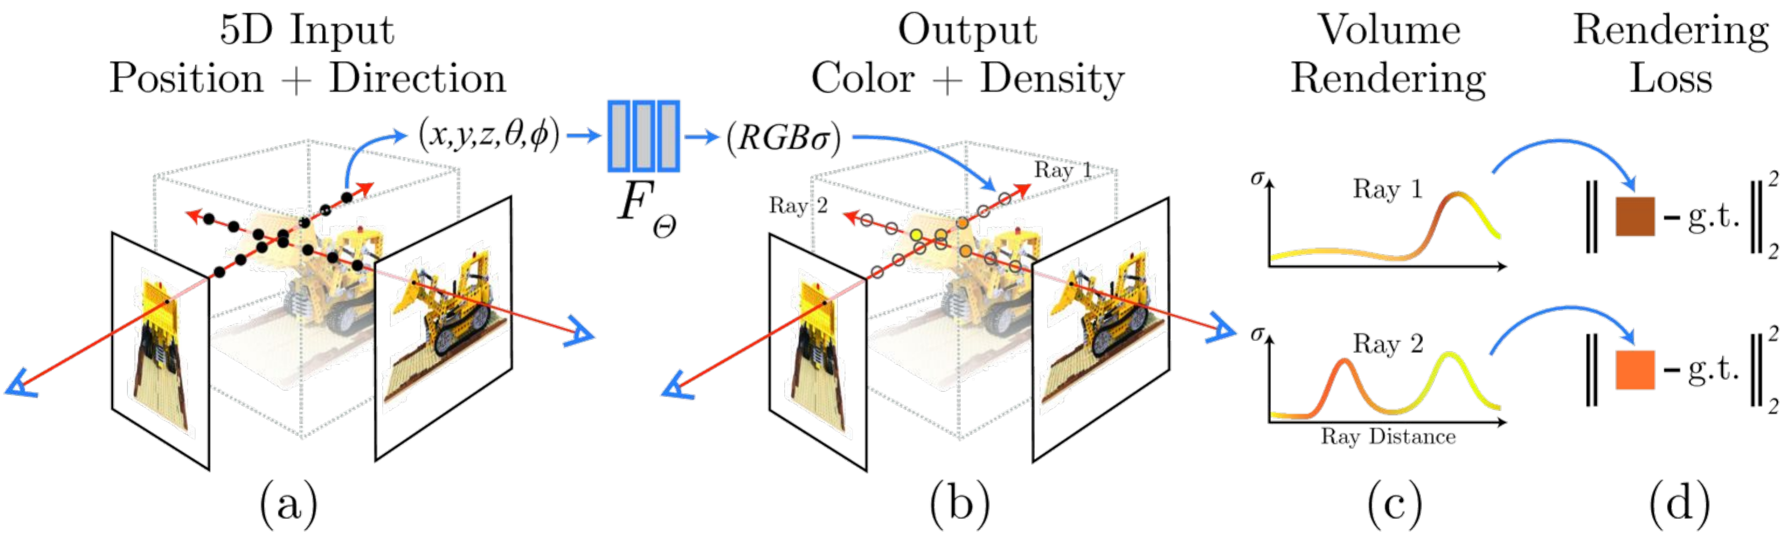

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Differentiable Volume Rendering

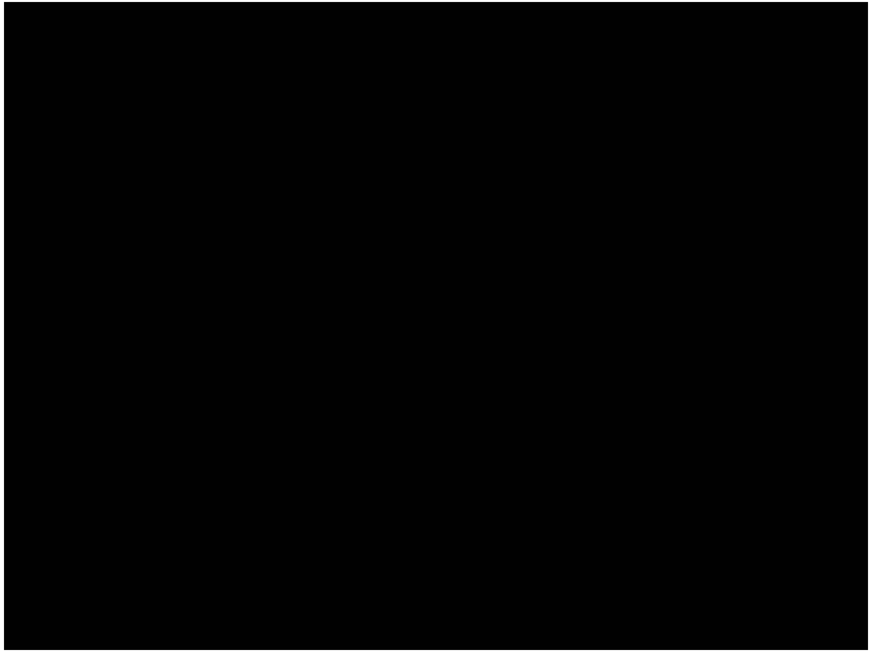

view synthesis

UNIVERSITY OF CAMBRIDGE

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Differentiable Volume Rendering

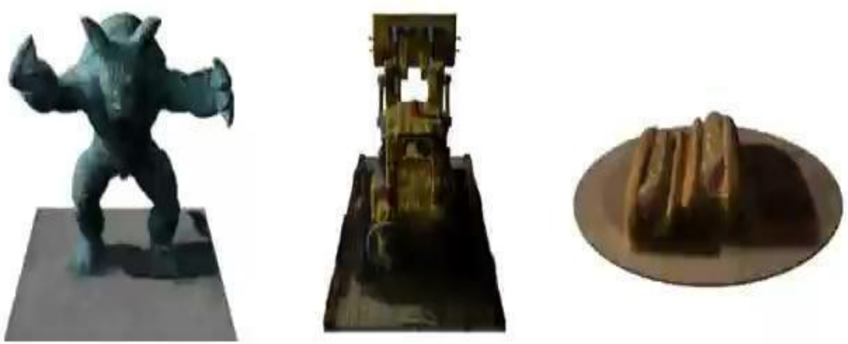

relighting

Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B. and Barron, J.T. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2021 (pp. 7495-7504).

Differentiable Volume Rendering

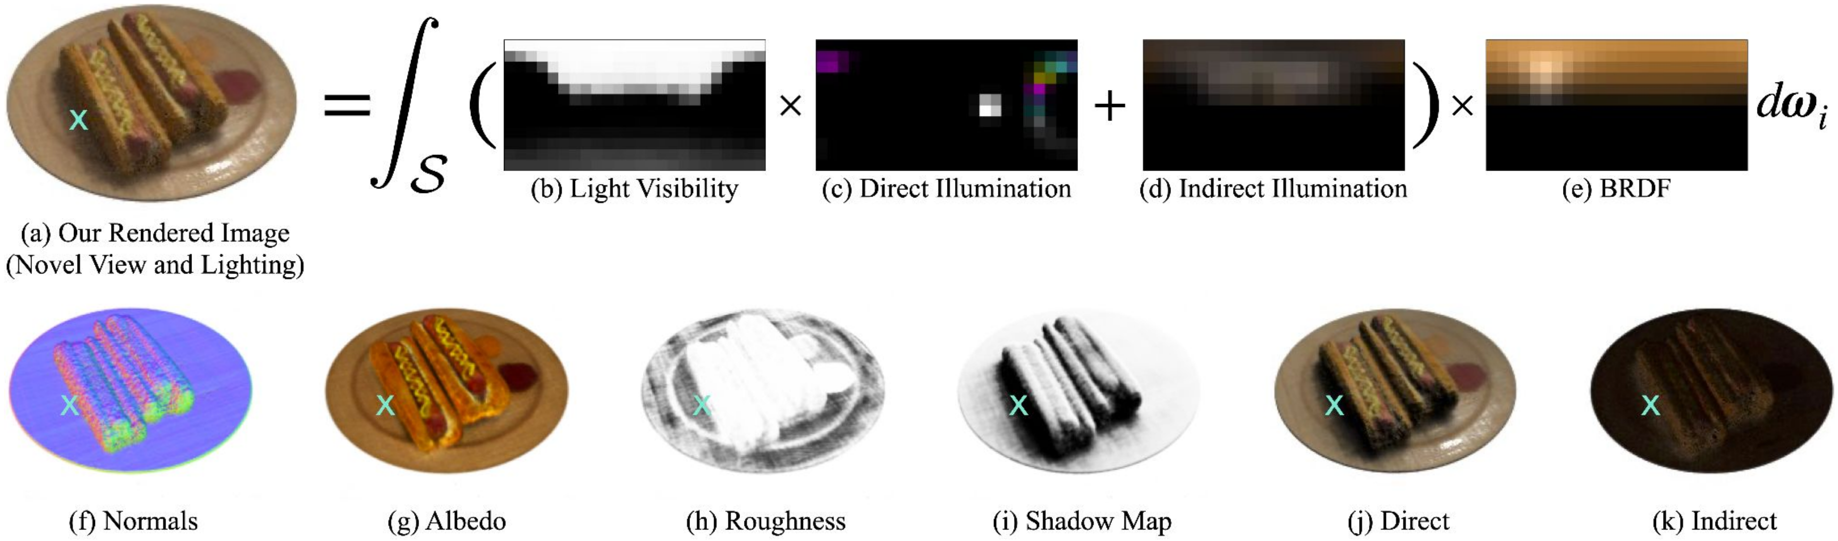

Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B. and Barron, J.T. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2021 (pp. 7495-7504).

Differentiable Volume Rendering

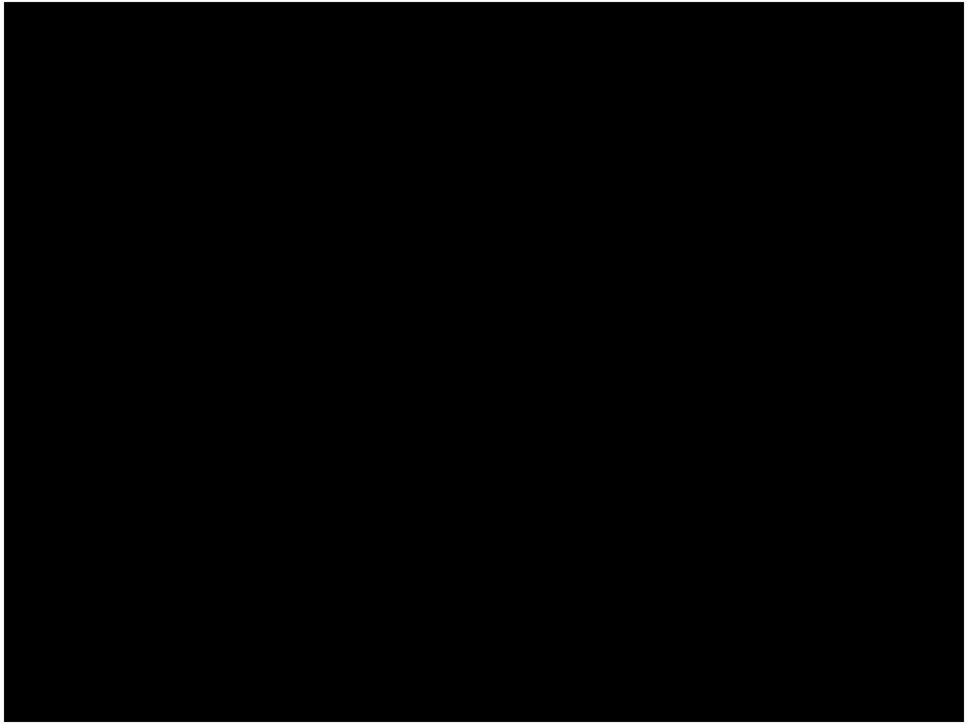

Zhang, Jason, et al. "NeRS: neural reflectance surfaces for sparse-view 3D reconstruction in the wild." *Advances in Neural Information Processing Systems* 34 (2021): 29835-29847.

Differentiable Volume Rendering

UNIVERSITY OF CAMBRIDGE

Tianhao Wu, Fangcheng Zhong, Andrea Tagliasacchi, Forrester Cole, and Cengiz Oztireli. D2NeRF: Self-Supervised Decoupling of Dynamic and Static Objects from a Monocular Video. In NeurIPS, 2022

Differentiable Graphics

- Unified framework to simultaneously infer multiple scene parameters
- Self-supervision
- Generalisability
- Cross regularisation
- Physics consistency in geometry and lighting