Differentiable 3D Visual Computing

Fangcheng Zhong

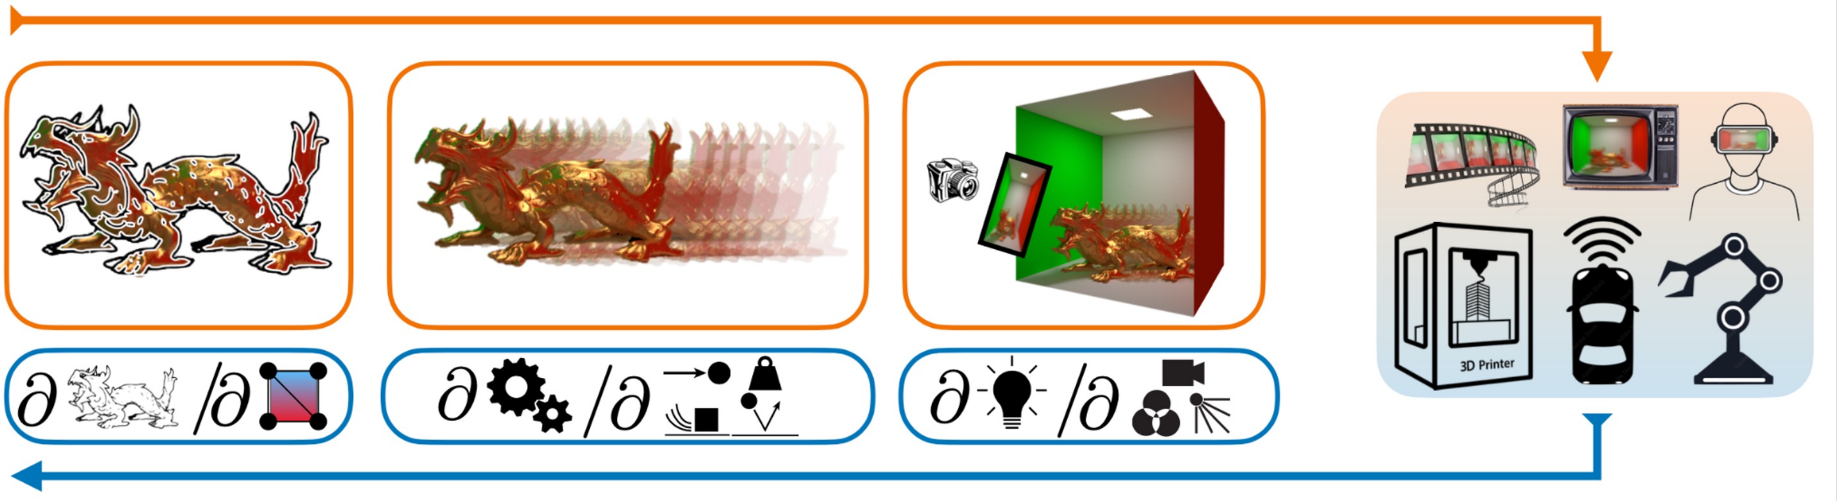

UNIVERSITY OF
CAMBRIDGE

Visual computing

Visual computing refers to any algorithm with visual content being the input or output

Machine visual perception

Computer graphics

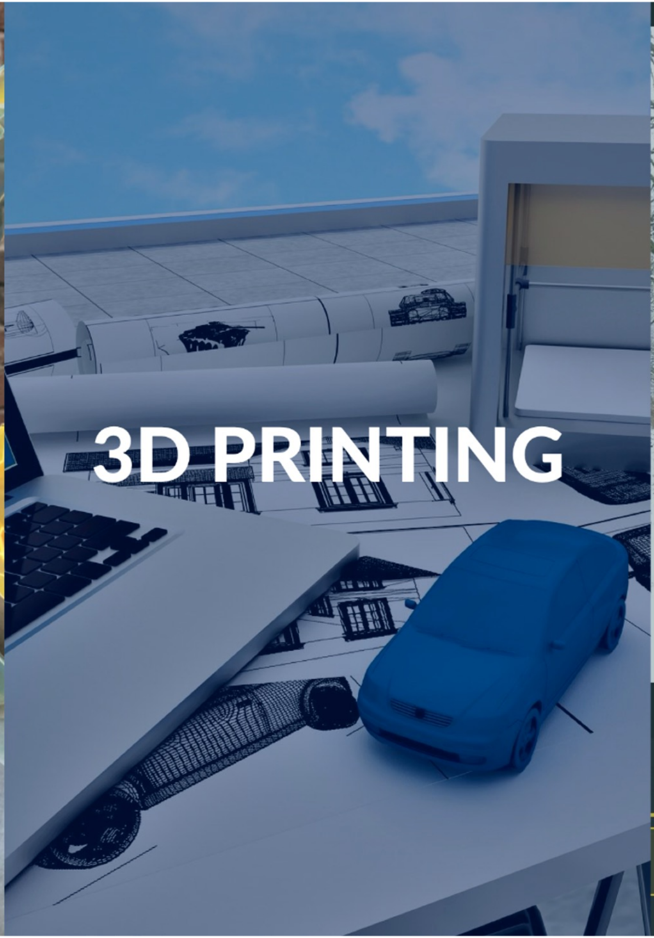

UNIVERSITY OF
CAMBRIDGE

3D visual computing

Any algorithm with 3D visual content being the input or output, a bridge between reality and digital reality!

3D visual computing

Real scene

3D representation

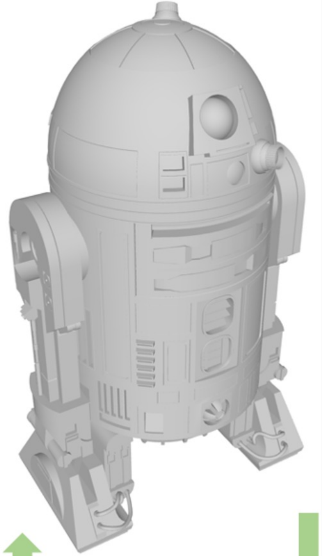

Virtual scene

Image-to-scene inference (computer vision)

Rendering (computer graphics)

3D visual computing

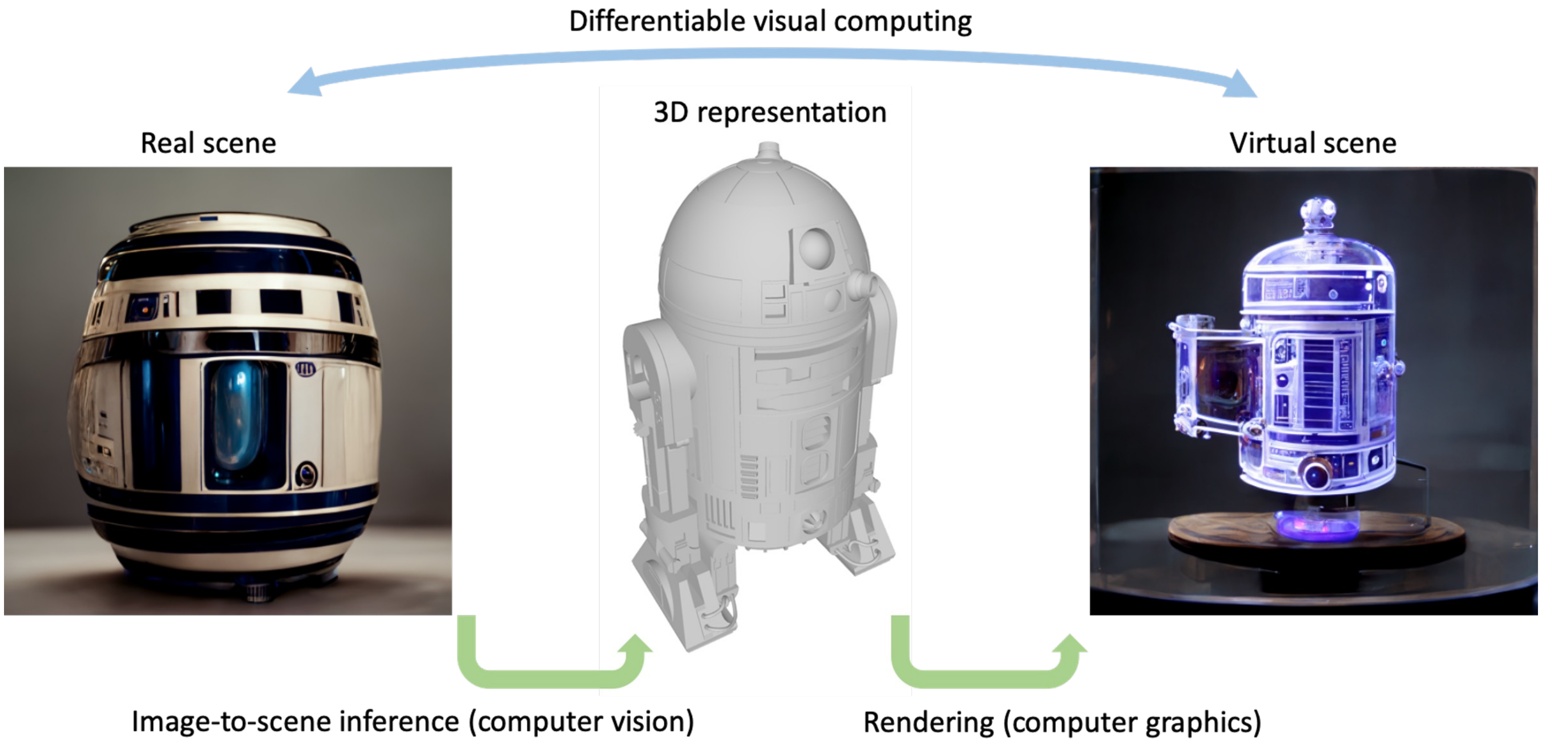

3D scene representation

Geometry

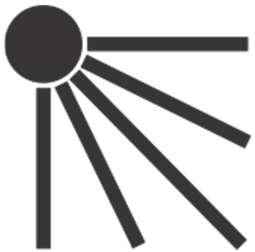

Lighting

Materials

Motion

Rendering

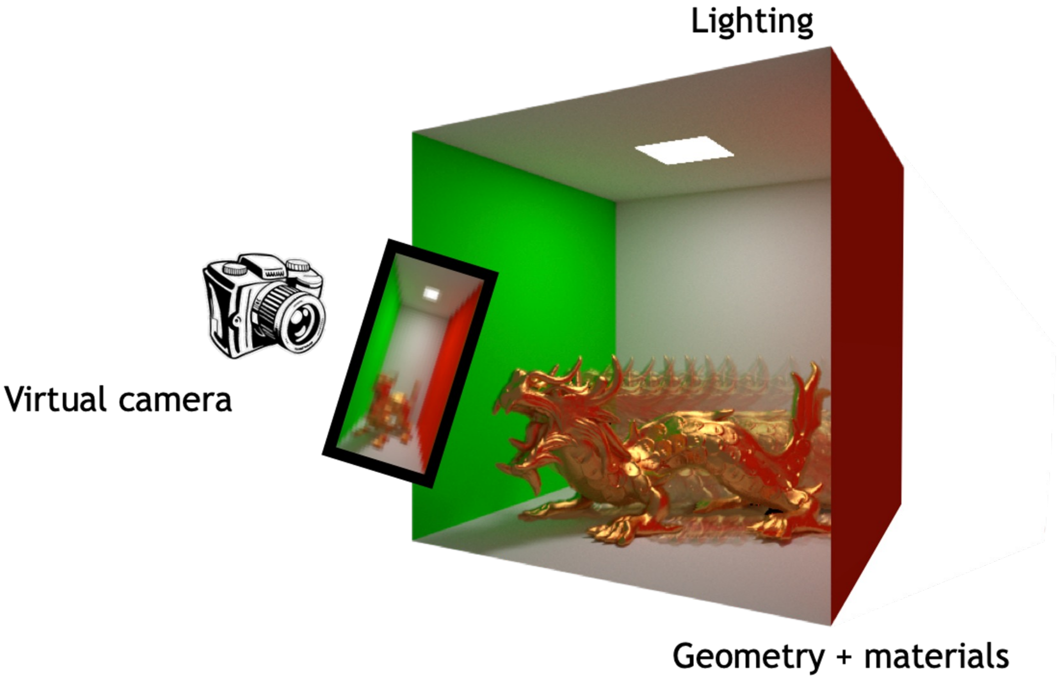

Image-to-scene inference

Traditional approach

- structured light
- multi-view stereo
- motion capture
- photometric stereo
- …

Image-to-scene inference

Deep learning approach

- learn a mapping from images to 3D scene parameters
- use a large dataset with correspondence of images and scene parameters (similar to a regression problem!)

Image-to-scene inference

Inverse rendering approach

$$\underset{\mathbf{s}}{\operatorname{argmin}} \sum_{i,t} \| R(\mathbf{s}, \mathbf{c}_{i,t}) - I_{i,t} \|$$

s *scene parameters, could be a function time t*

R *rendering operator*

c *camera parameters at the i-th view and time t*

I *image at the i-th view and time t*

Inverse rendering

$$\underset{\mathbf{s}}{\operatorname{argmin}} \sum_{i,t} \| R(\mathbf{s}, \mathbf{c}_{i,t}) - I_{i,t} \|$$

R is not differentiable in traditional graphics!

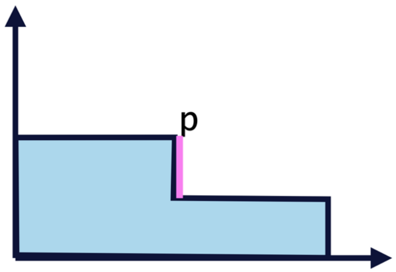

Differentiable rendering

$$\operatorname*{argmin}_{\mathbf{s}} \sum_{i,t} \| R(\mathbf{s}, \mathbf{c}_{i,t}) - I_{i,t} \|$$

Make rendering differentiable!

Differentiable rendering

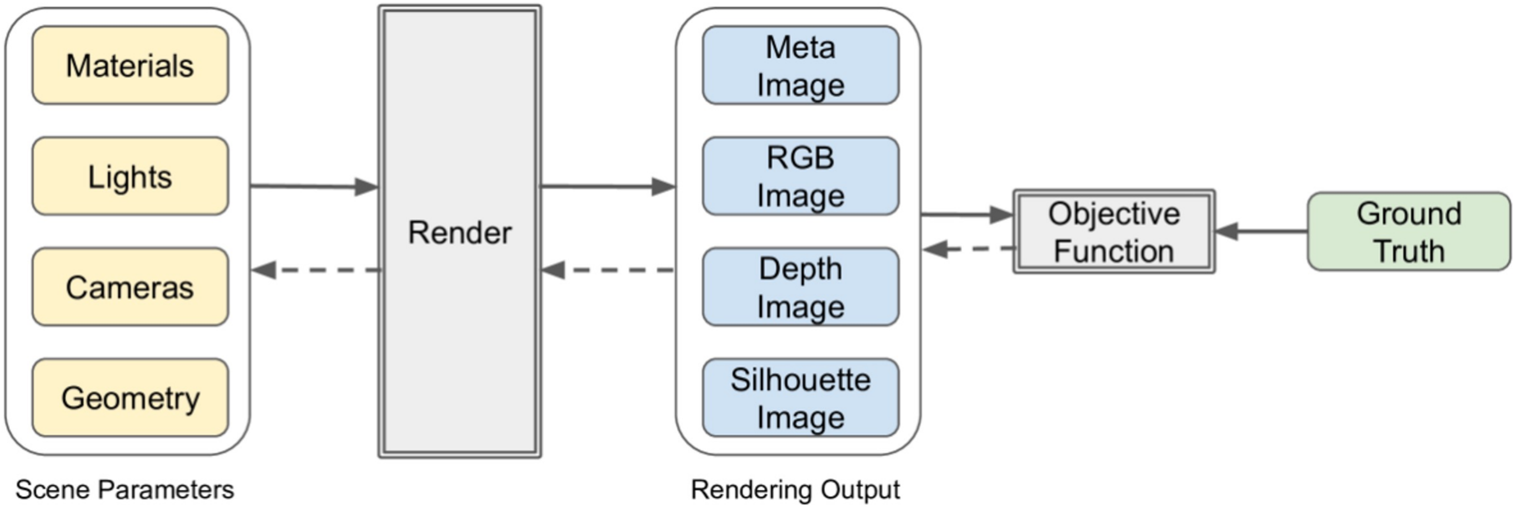

Optimisation with a differentiable renderer

Differentiable rendering

Make rendering differentiable!

- self-supervision
- generalisable to all scenarios
- consistency in geometry and light transport
- unified framework to simultaneously infer multiple scene parameters
- applications in physical inference, optimal control, scene understanding, computational design, manufacturing, autonomous vehicles, and robotics

Differentiable rendering

Make rendering differentiable!

- inference not in real time
- choice renderer
- initialisation

SoftRas: differentiable rasterization

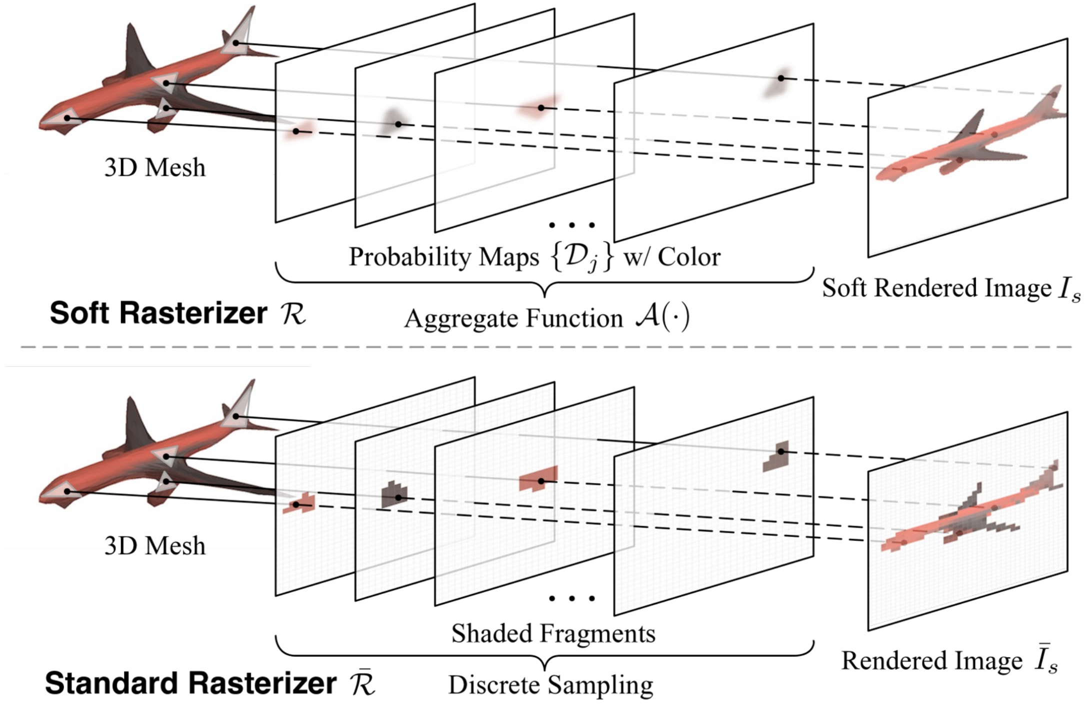

Liu, Shichen, Tianye Li, Weikai Chen, and Hao Li. "Soft rasterizer: A differentiable renderer for image-based 3d reasoning." In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7708-7717. 2019.

Application in surface reconstruction

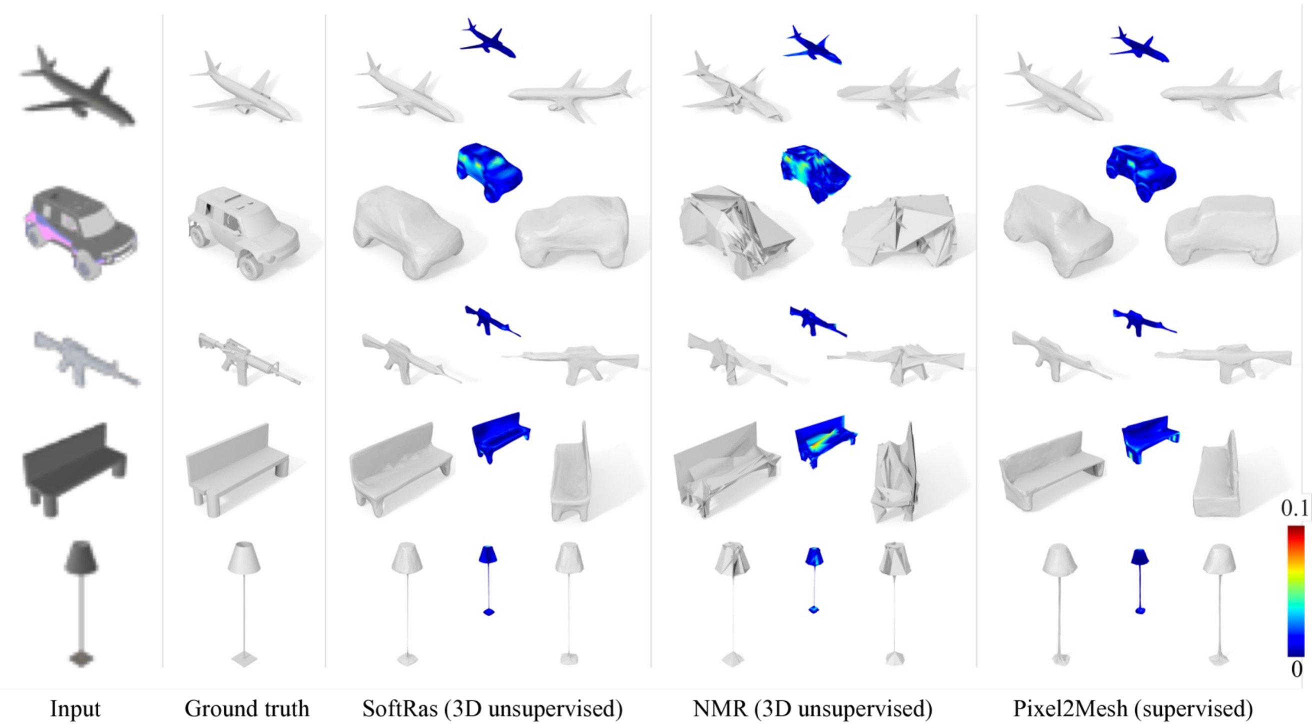

UNIVERSITY OF CAMBRIDGE

Liu, Shichen, Tianye Li, Weikai Chen, and Hao Li. "Soft rasterizer: A differentiable renderer for image-based 3d reasoning." In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7708-7717. 2019.

Application in extended reality

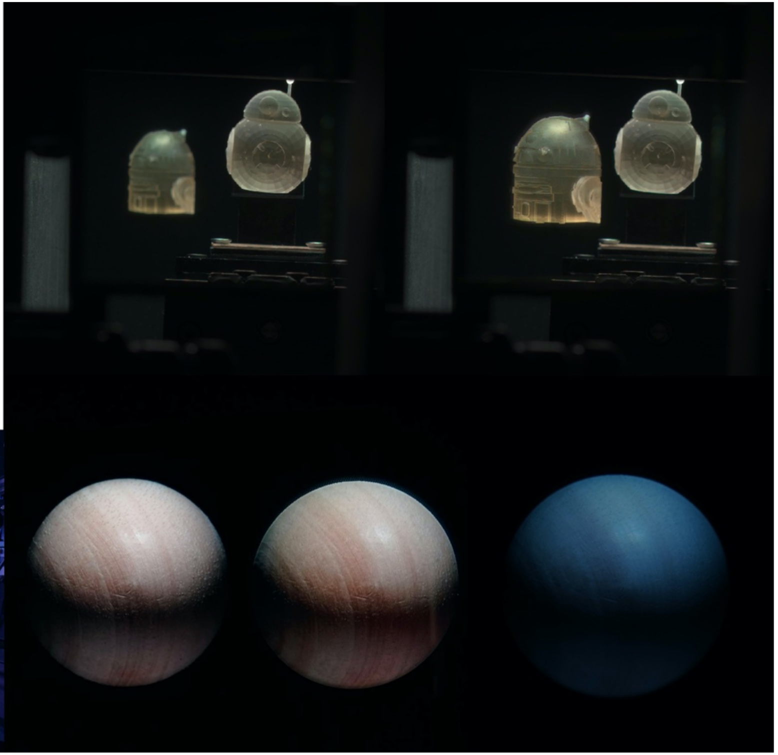

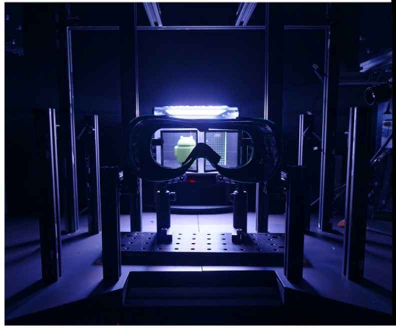

Photographs of physical objects next to virtual 3D objects rendered by a 3D display

UNIVERSITY OF CAMBRIDGE

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 Trans. Graph. 40, 6, Article 241 (December 2021) https://doi.org/10.1145/3478513.3480513

NeRF: differentiable volume rendering

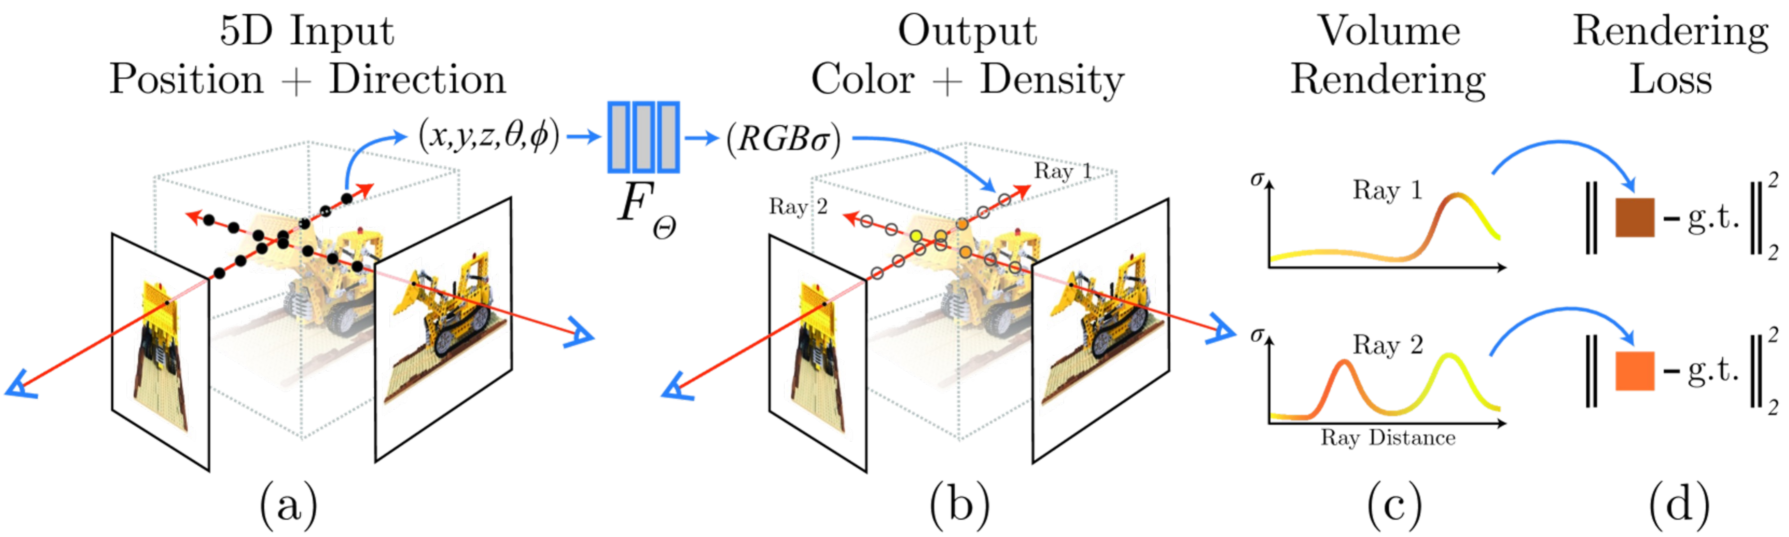

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020

Application in 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

Application in inferring illumination and materials

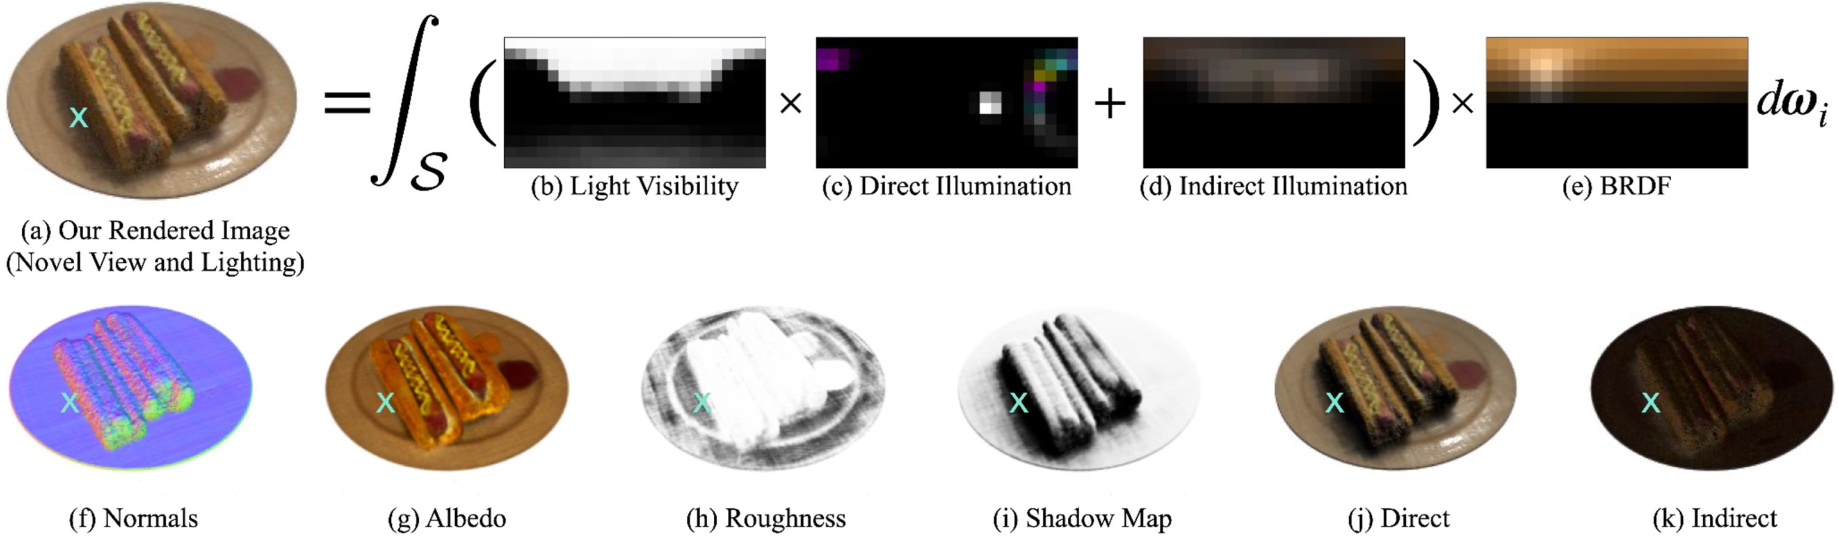

UNIVERSITY OF CAMBRIDGE

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).

Application in inferring illumination and materials

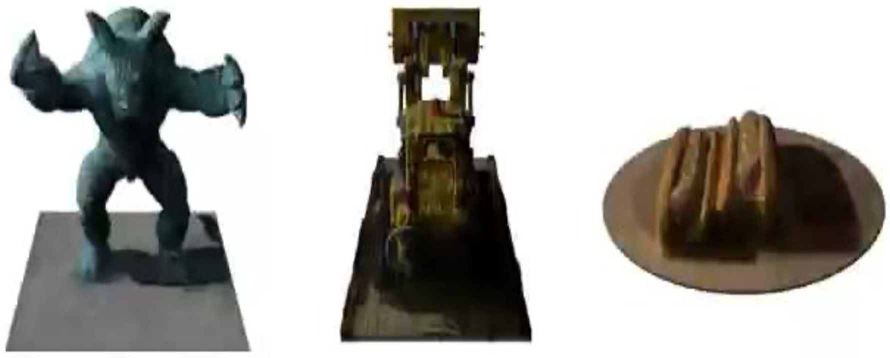

Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B. and Barron, J.T. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2021 (pp. 7495-7504).

Application in decoupling motion

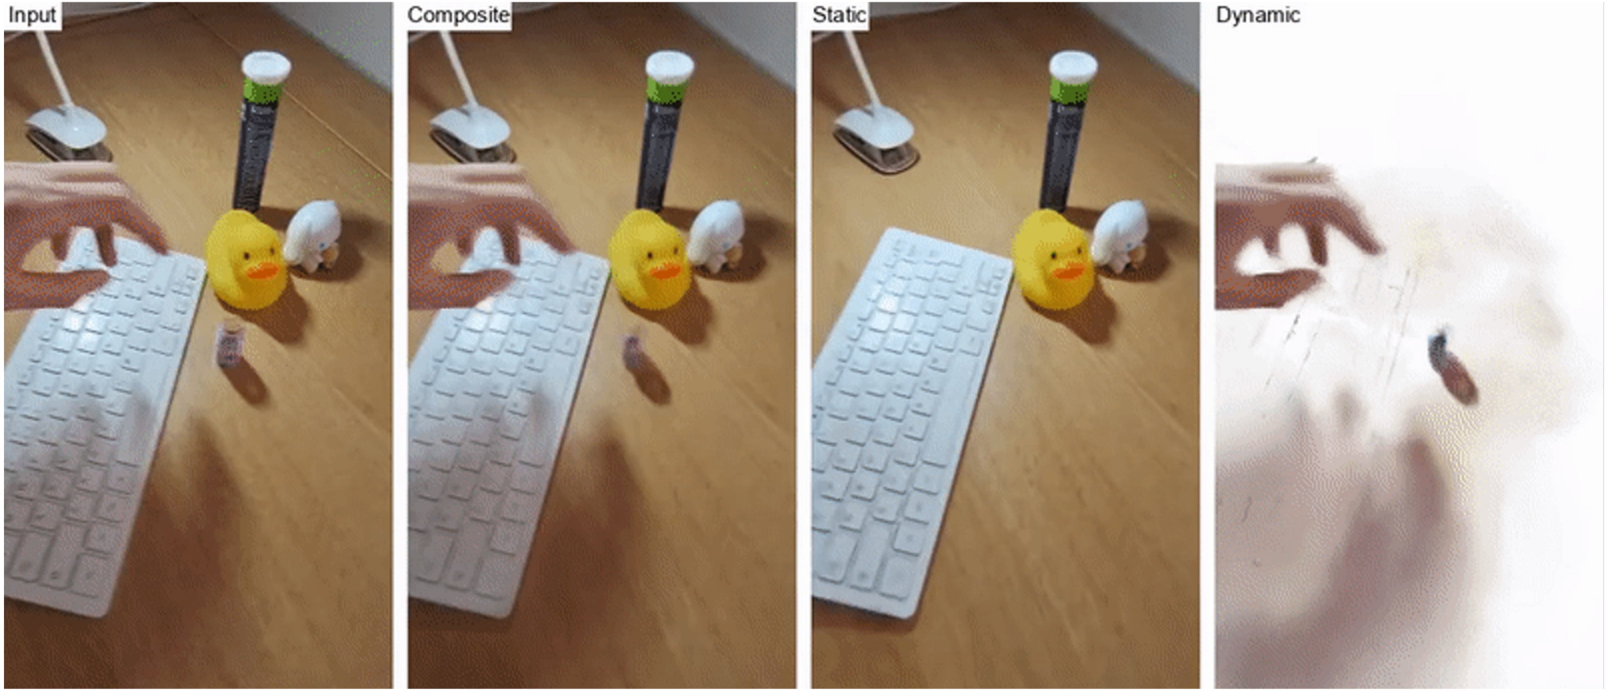

Tianhao Wu, Fangcheng Zhong, Andrea Tagliasacchi, Forrester Cole, and Cengiz Oztireli. D2NeRF: Self-Supervised Decoupling of Dynamic and Static Objects from a Monocular Video. In NeurIPS, 2022

Combining differentiable rendering with 3D generative modelling

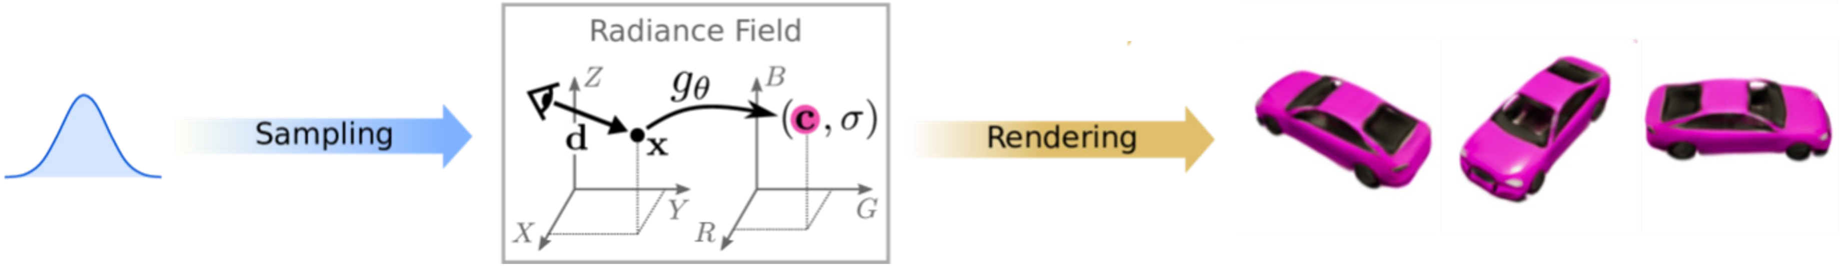

Schwarz, Katja, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. "Graf: Generative radiance fields for 3d-aware image synthesis." *Advances in Neural Information Processing Systems* 33 (2020): 20154-20166

Combining differentiable rendering with generative modelling

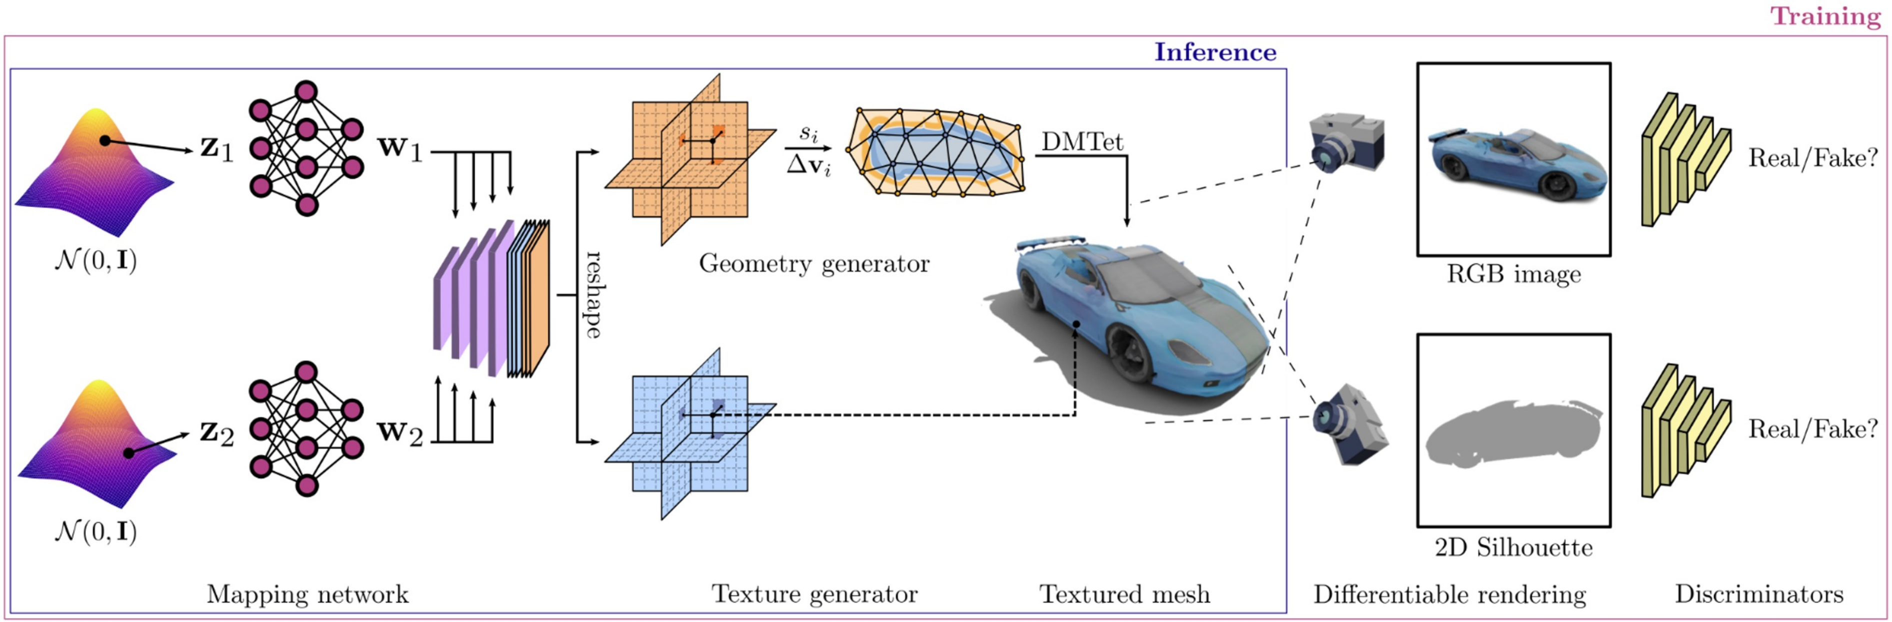

Gao, Jun, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. "GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images." *In NeurIPS, 2022*

Combining differentiable rendering with generative modelling

UNIVERSITY OF CAMBRIDGE

Gao, Jun, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. "GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images." *In NeurIPS, 2022*

DVC Frameworks

https://pytorch3d.org/

https://www.tensorflow.org/graphics

https://github.com/NVlabs/nvdiffrast

https://www.mitsuba-renderer.org/