Neural Rendering

Chuan Li

Lambda Labs

Collaborators: Thu Nguyen-Phuoc, Bing Xu, Yongliang Yang, Stephen Balaban, Lucas Theis, Christian Richardt, Junfei Zhang, Rui Wang, Kun Xu, Rui Tang

Lambda

Forward (Computer Graphics)

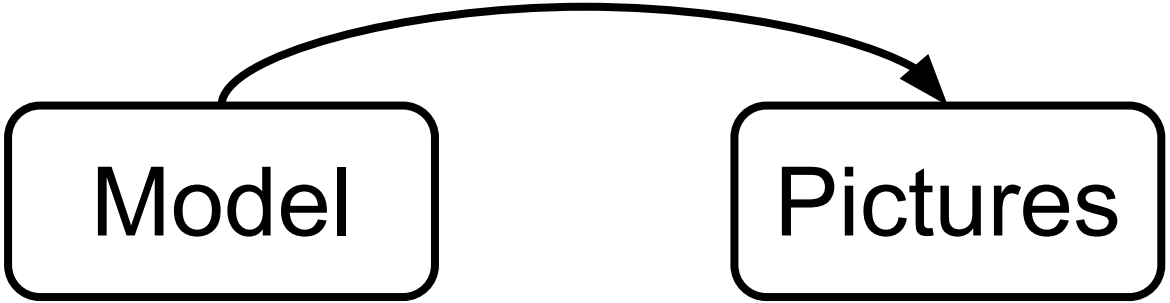

Forward (Computer Graphics)

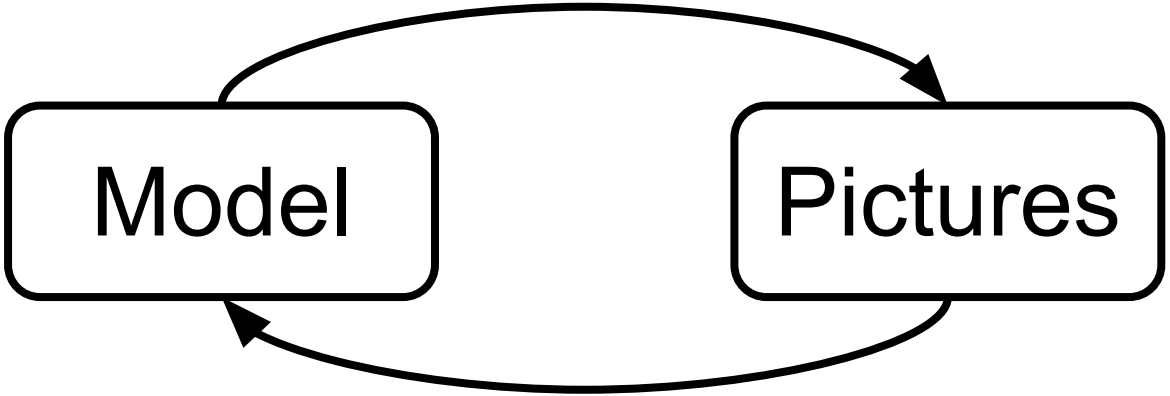

Inverse (Computer Vision)

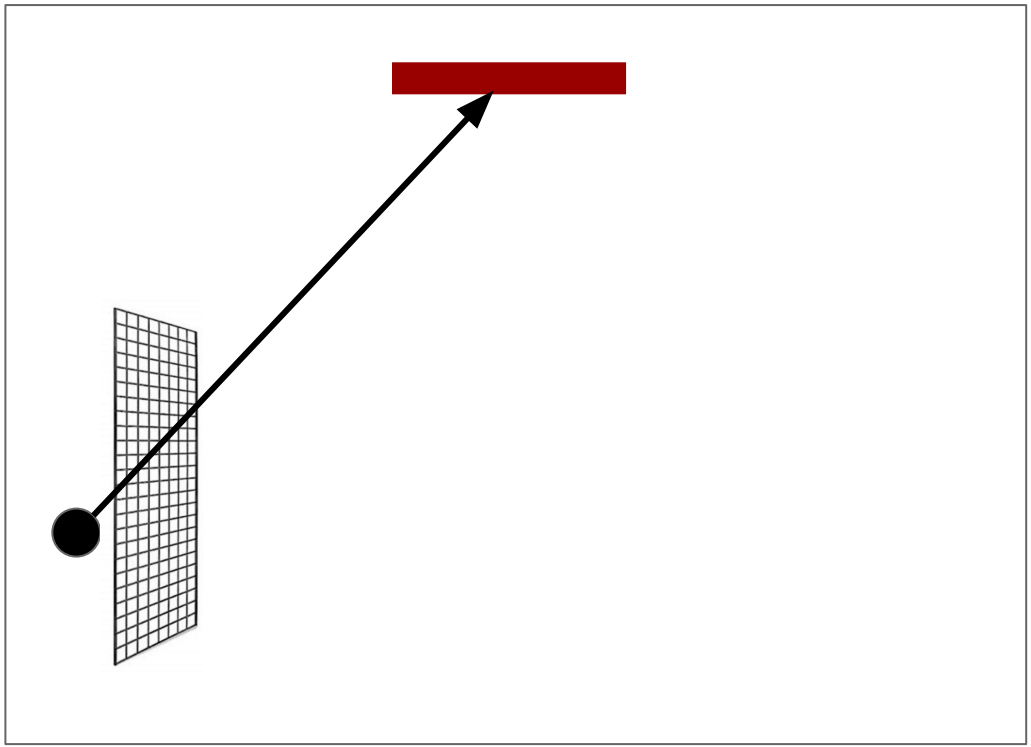

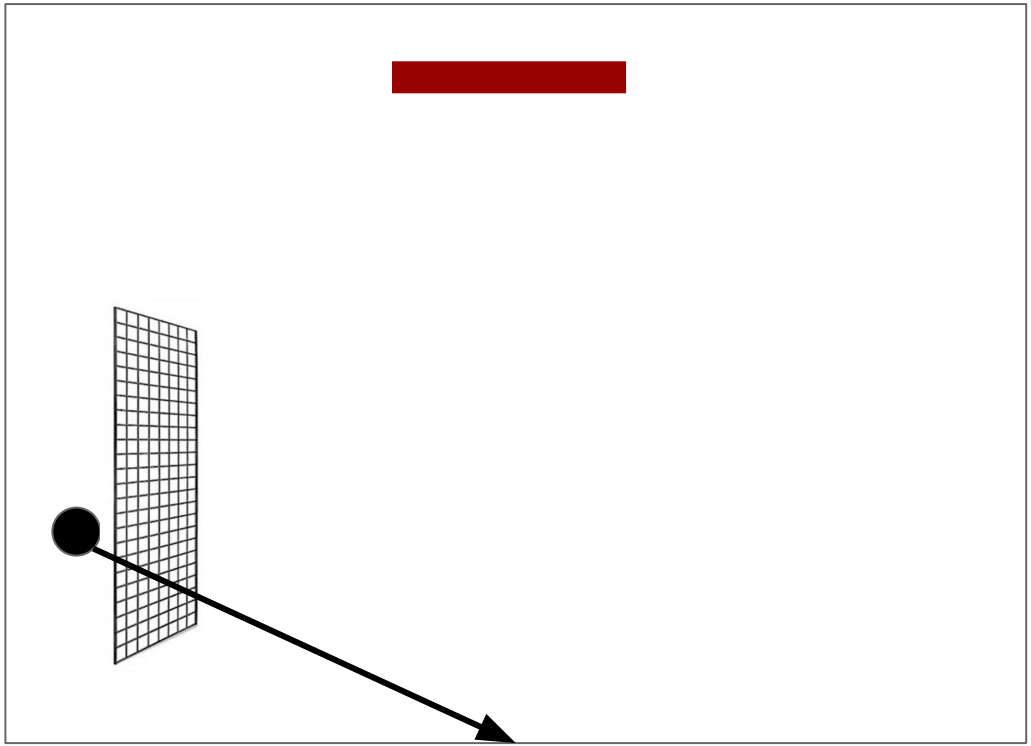

Lambda

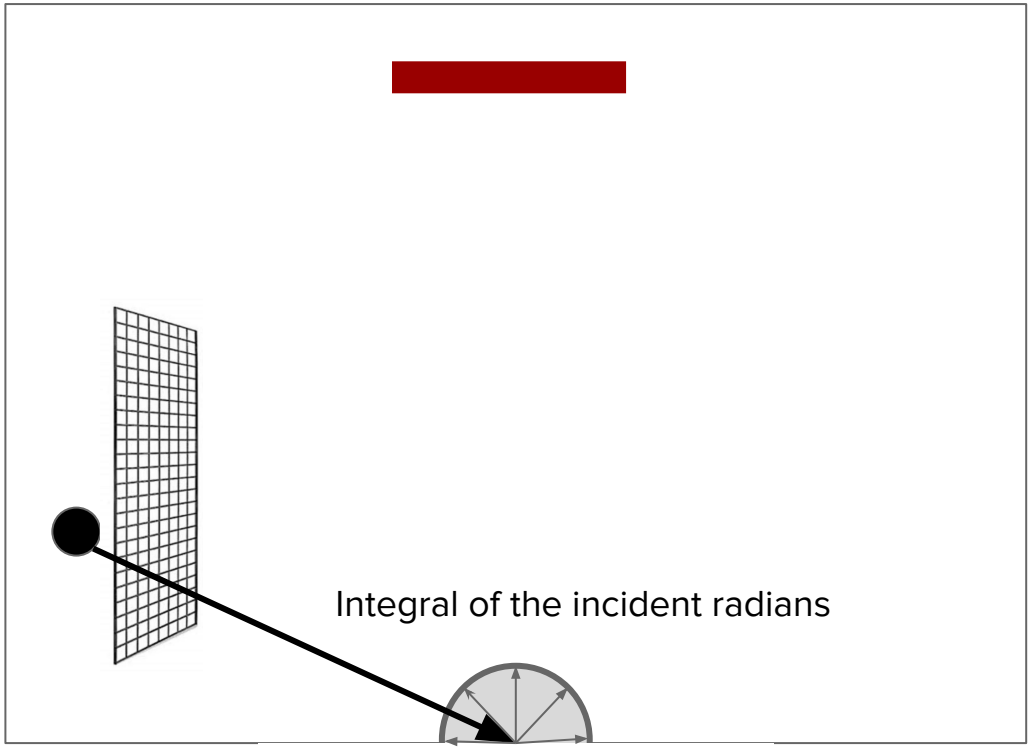

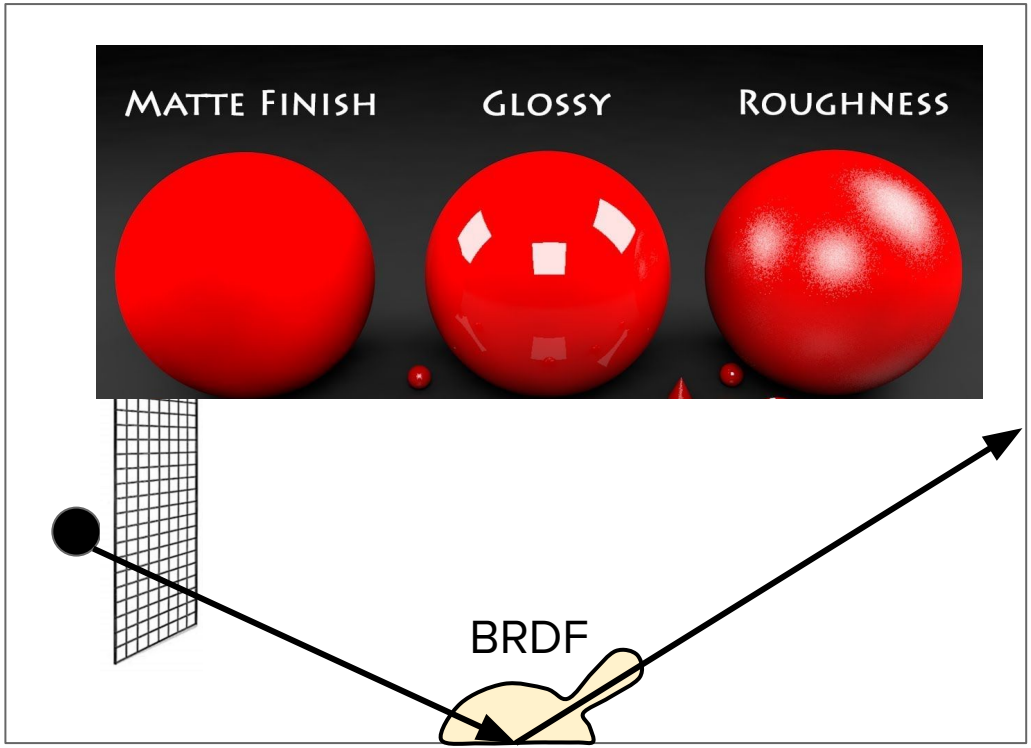

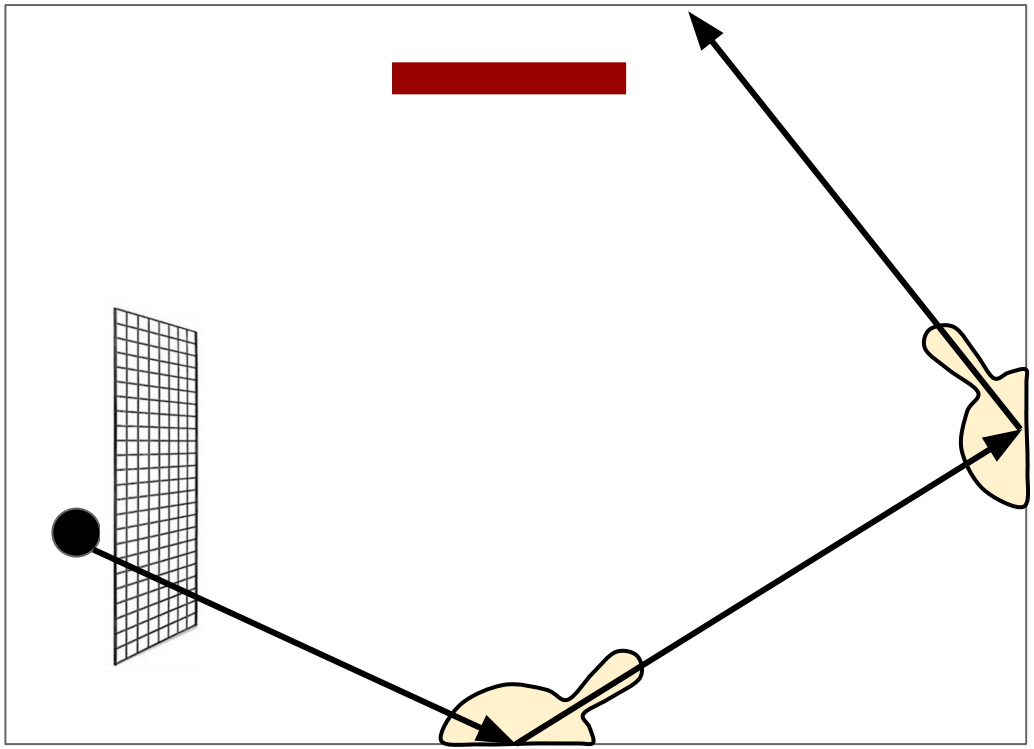

Next Event Estimation

Bidirectional Ray Tracing

Metropolis Light Transportation

Lambda

32K SPP Ray Tracing (90 mins 12 CPU Cores)
The Tungsten Renderer

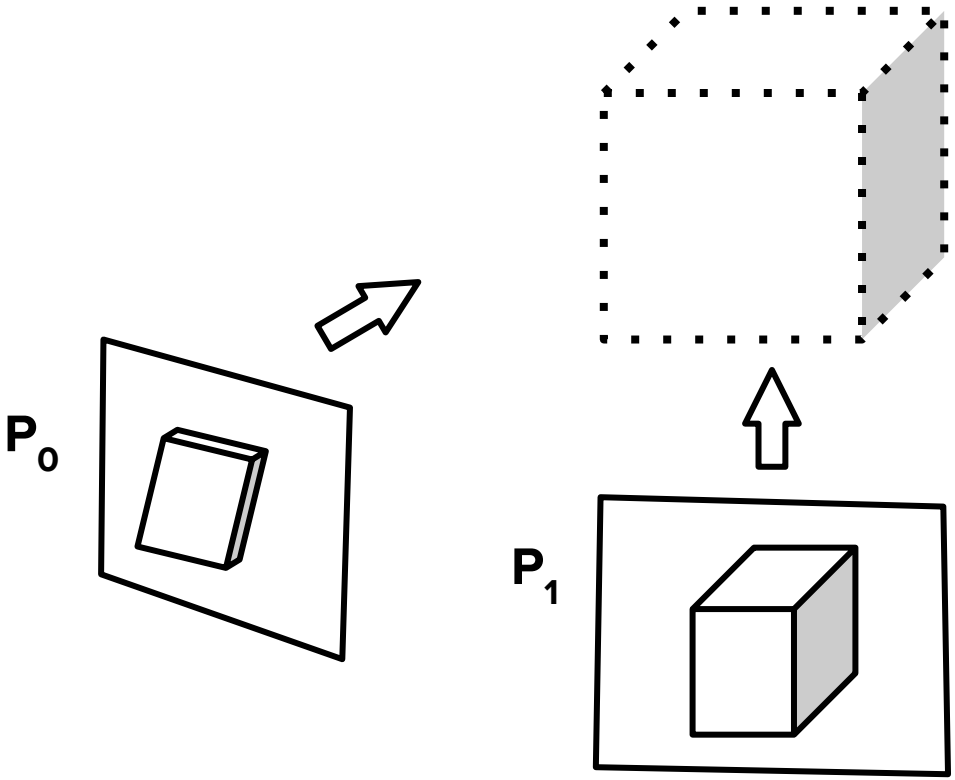

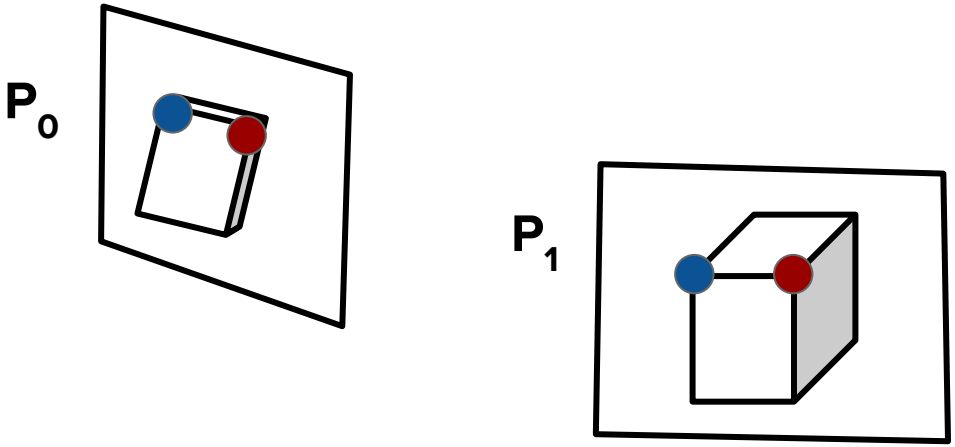

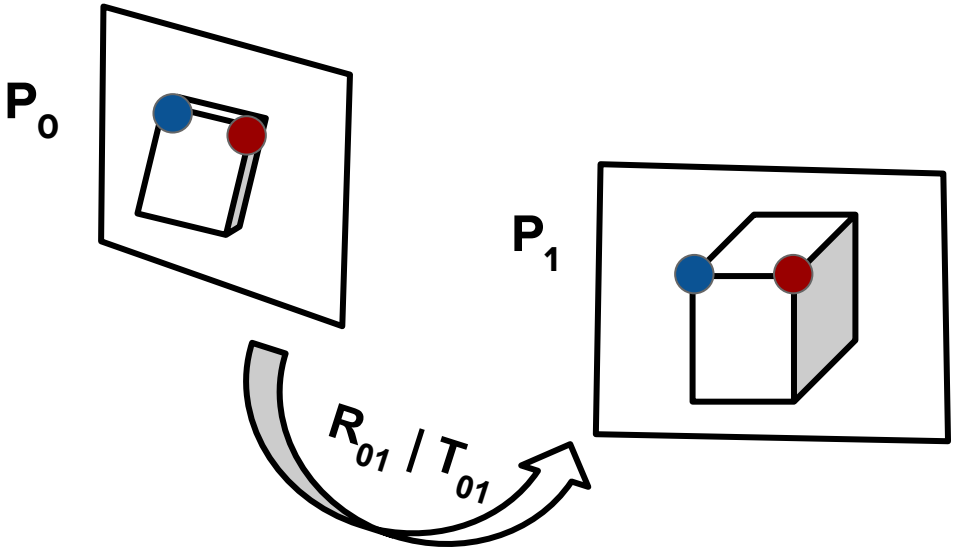

Lambda

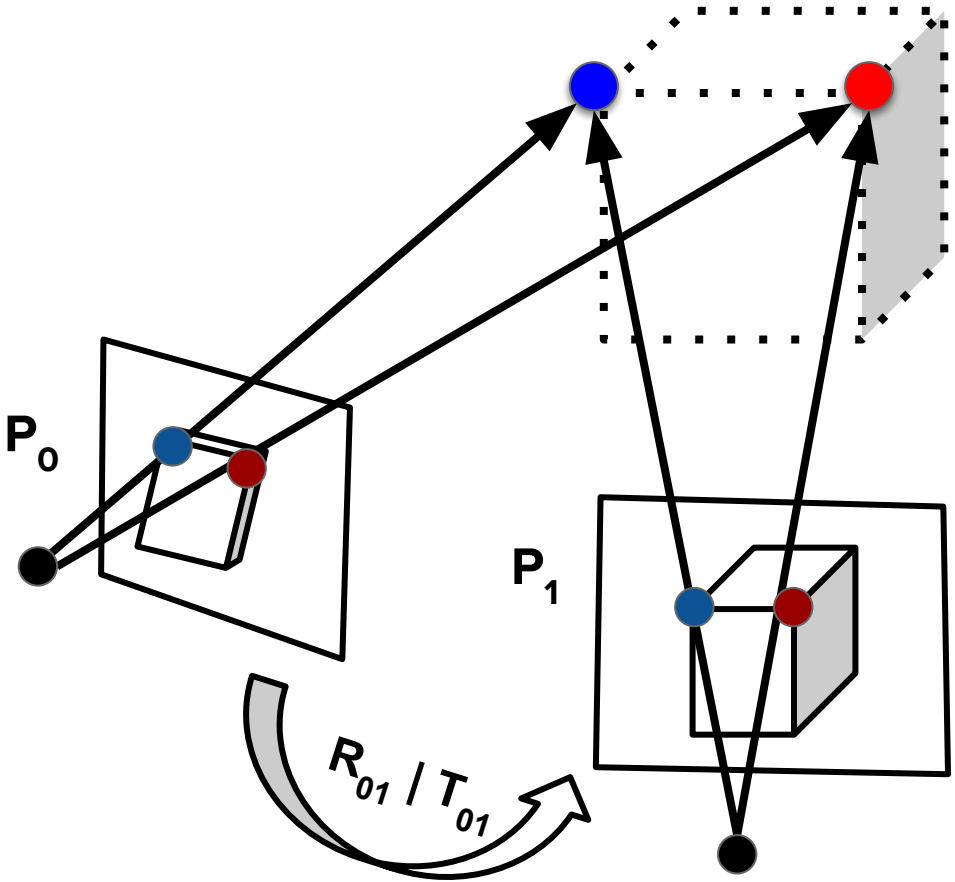

Lambda

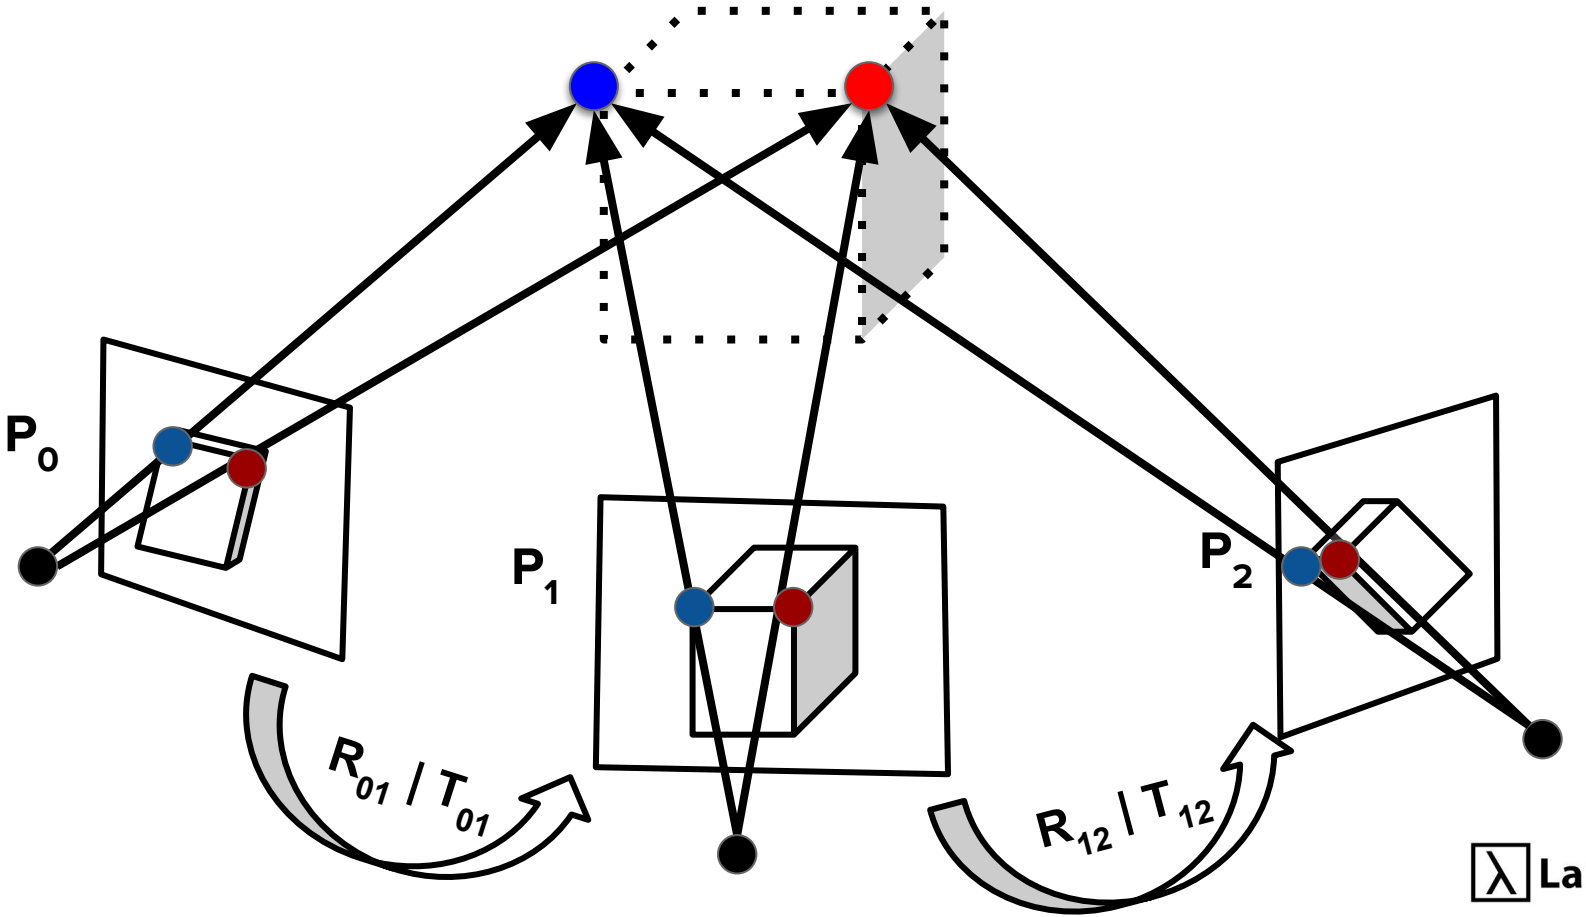
P0
P1
P2
R01 / T01
R12 / T12
Lambda

Building Rome in a Day

Sameer Agarwal, Noah Snavely, Ian Simon, Steven M. Seitz and Richard Szeliski

Lambda

Sub-module

End-2-End

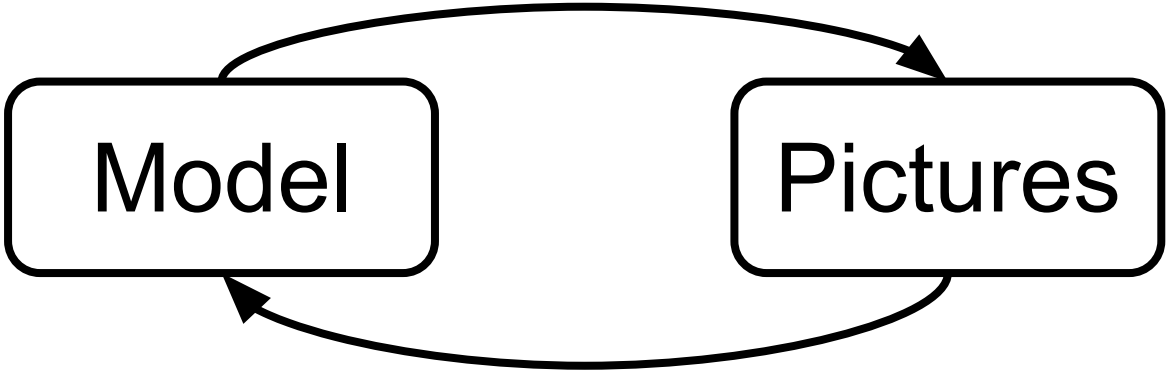

Differentiable Rendering

1 SPP

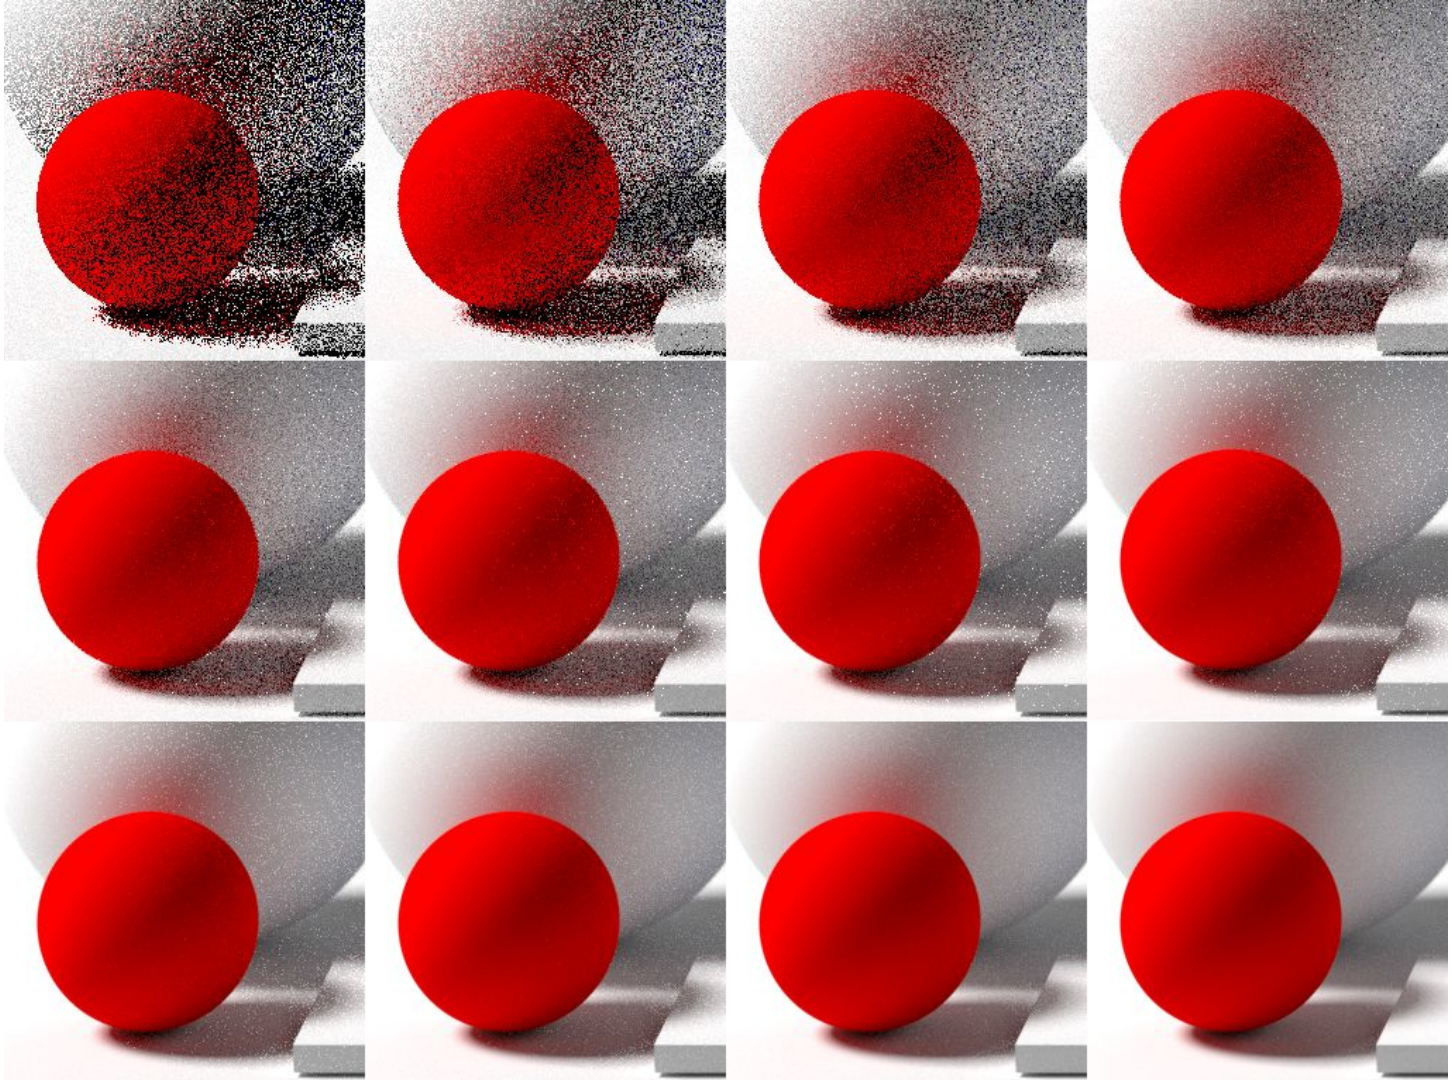

2048 SPP

Lambda

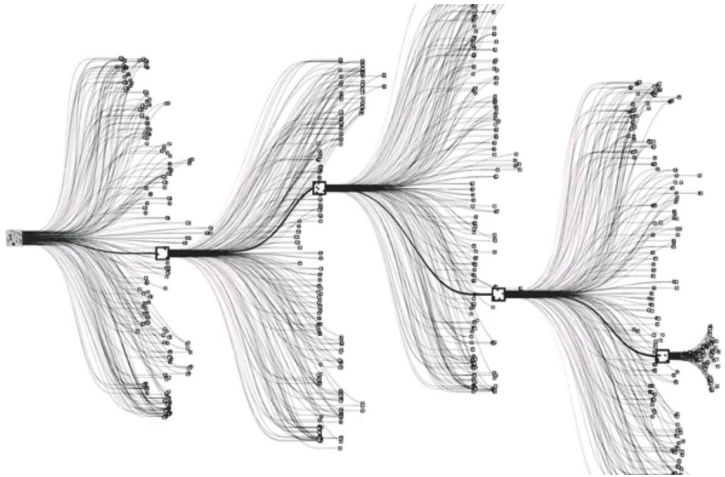

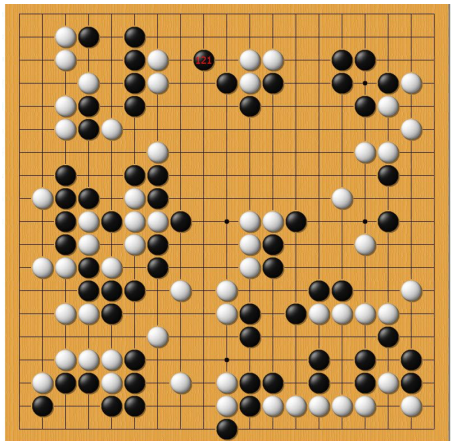

Mastering the game of Go with deep neural networks and tree search

David Silver et al.

Lambda

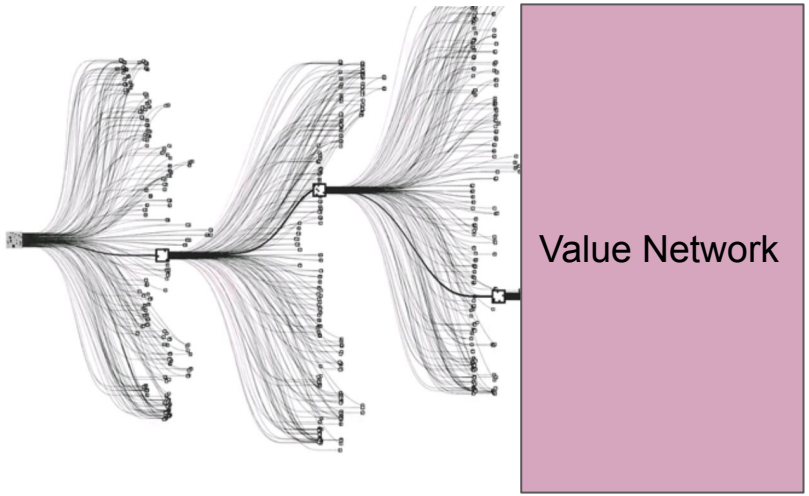

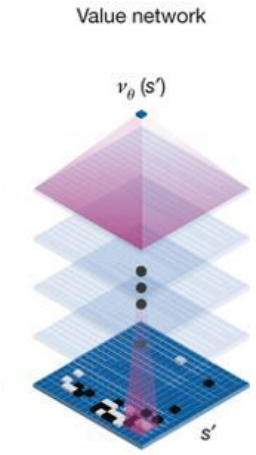

Mastering the game of Go with deep neural networks and tree search

David Silver et al.

Lambda

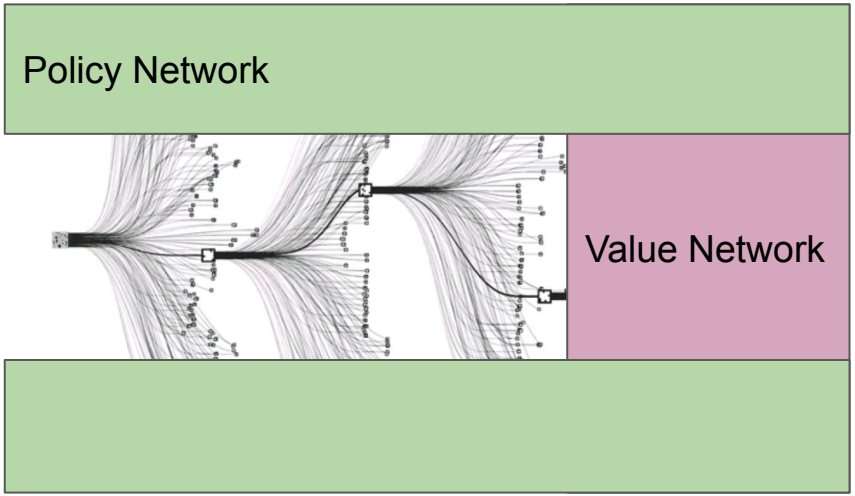

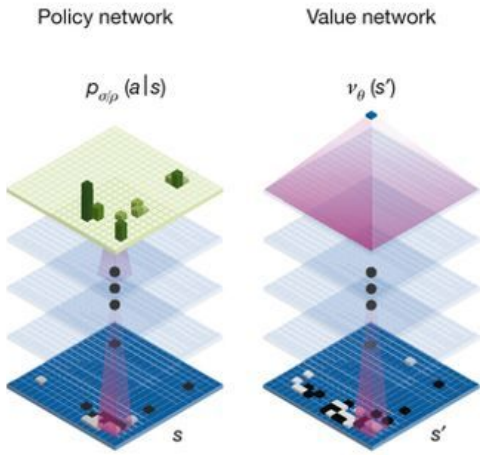

Mastering the game of Go with deep neural networks and tree search

David Silver et al.

Lambda

Value Networks

Denoising

4 SPP

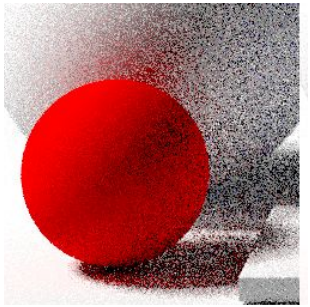

2^15 SPP

Lambda

Value Networks

Denoising

4 SPP

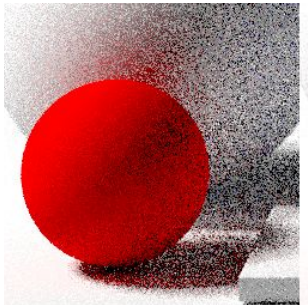

2^15 SPP

Policy Networks

Same SPP

Value Networks

Denoising

4 SPP

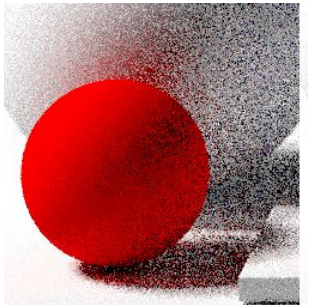

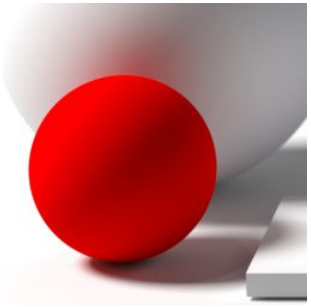

2^15 SPP

Policy Networks

Same SPP

4 SPP

Denoised
1 sec 2080 Ti

32K SPP Ray Tracing
90 mins 12 cores CPU

Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation

B Xu et al. Siggraph Asia 2019

Lambda

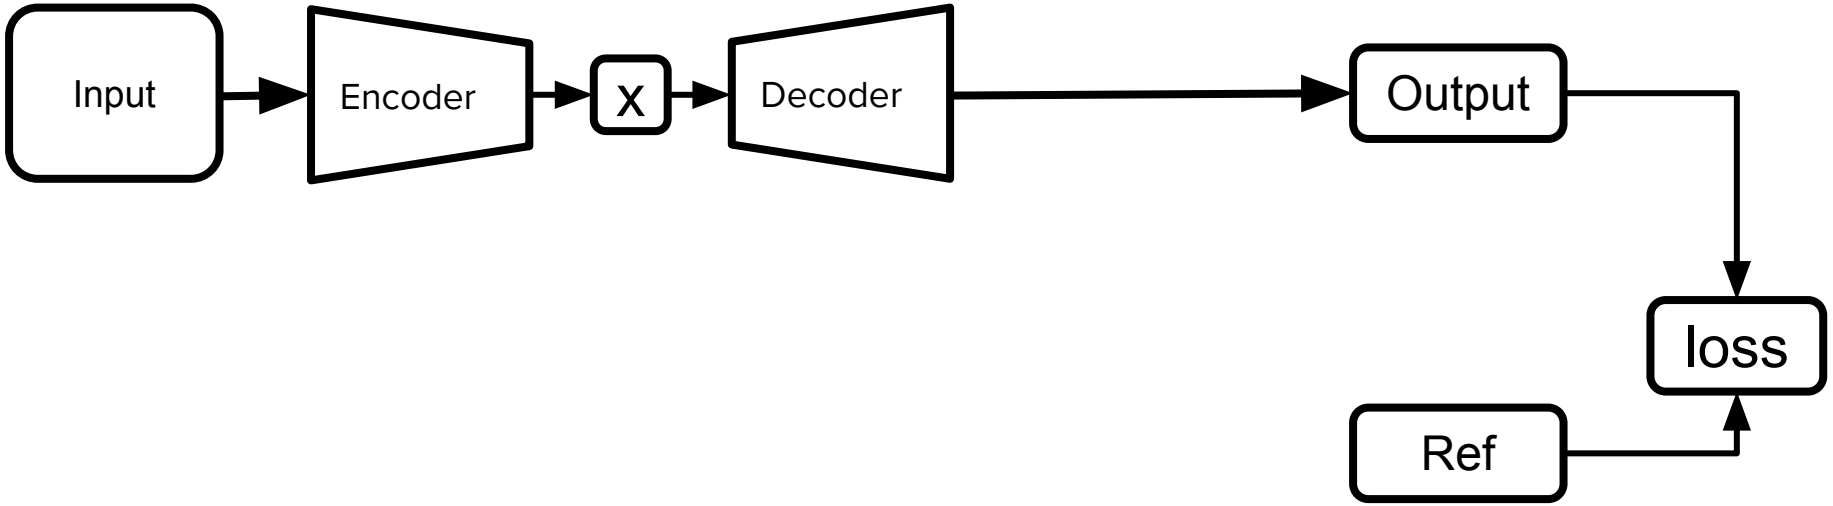

Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation

B Xu et al. Siggraph Asia 2019

L1 VGG Loss

L1 VGG Loss + GAN

Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation

B Xu et al. Siggraph Asia 2019

Lambda

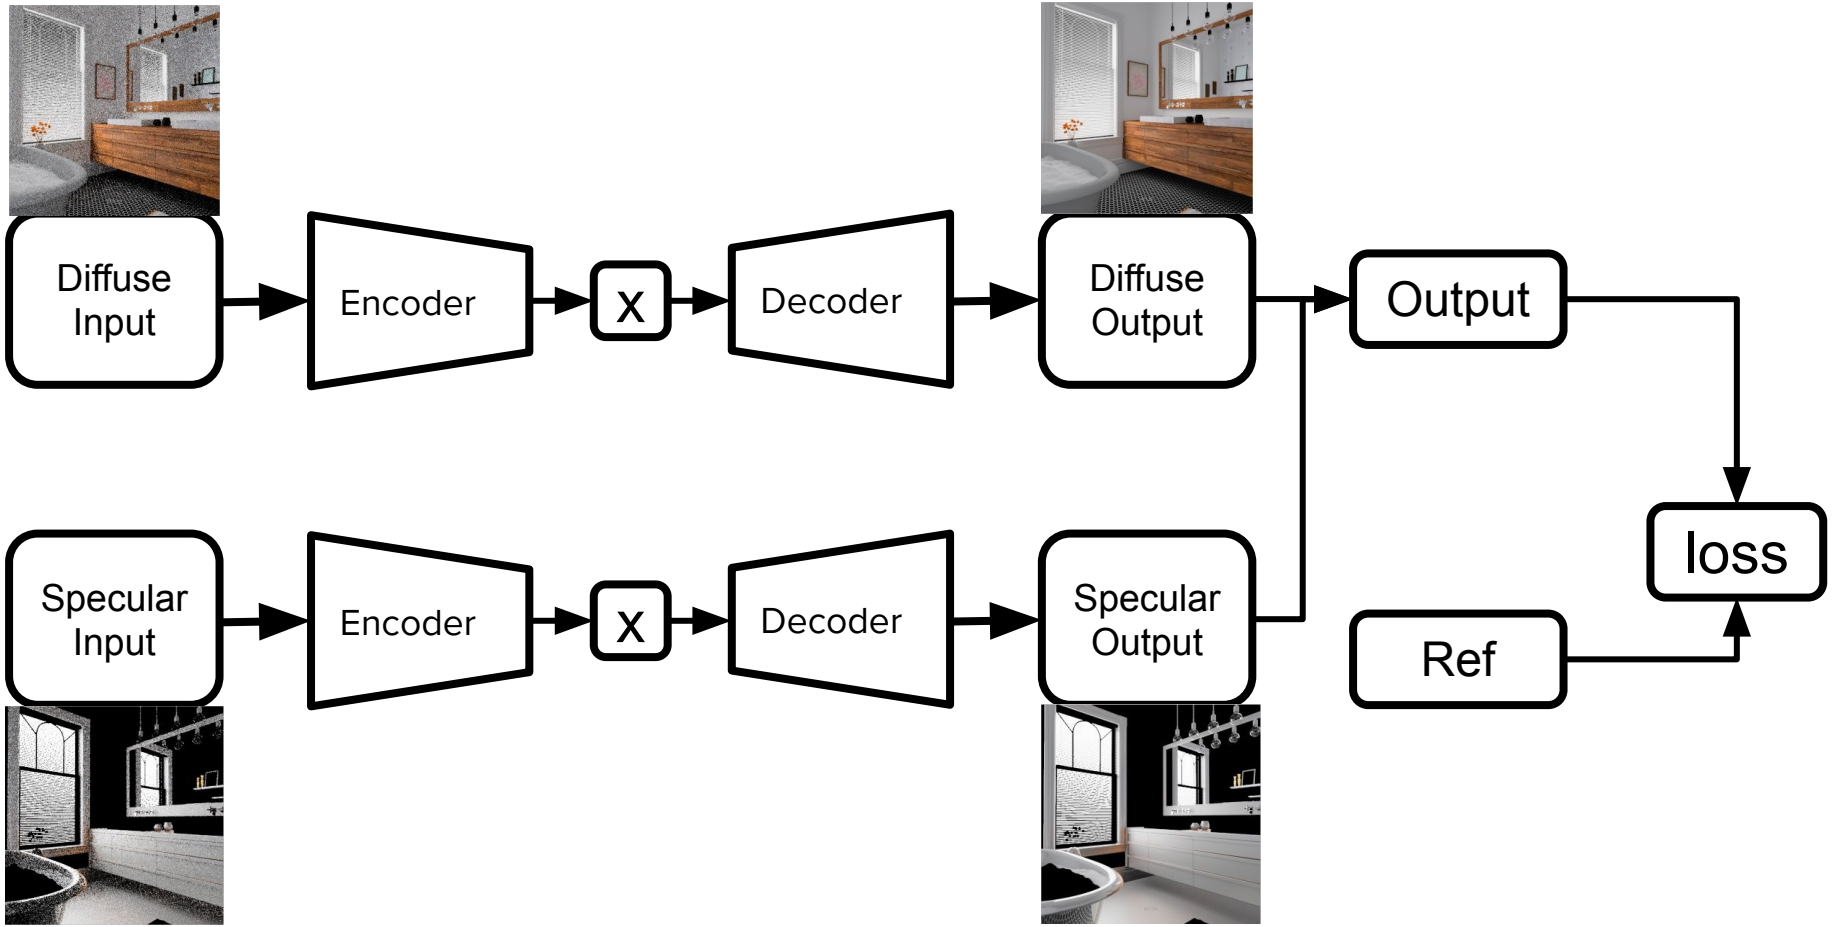

Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation

B Xu et al. Siggraph Asia 2019

Lambda

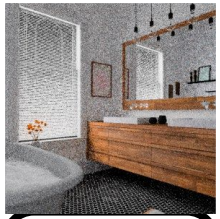

Diffuse Input

Encoder

X

Decoder

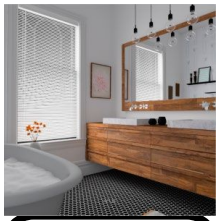

Diffuse Output

Output

loss

Specular Input

Encoder

X

Decoder

Specular Output

Ref

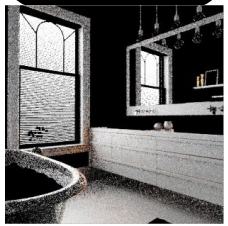

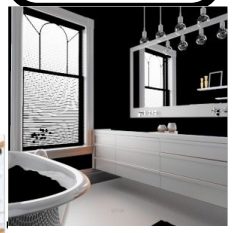

Auxiliary

Albedo, normal, depth

Lambda

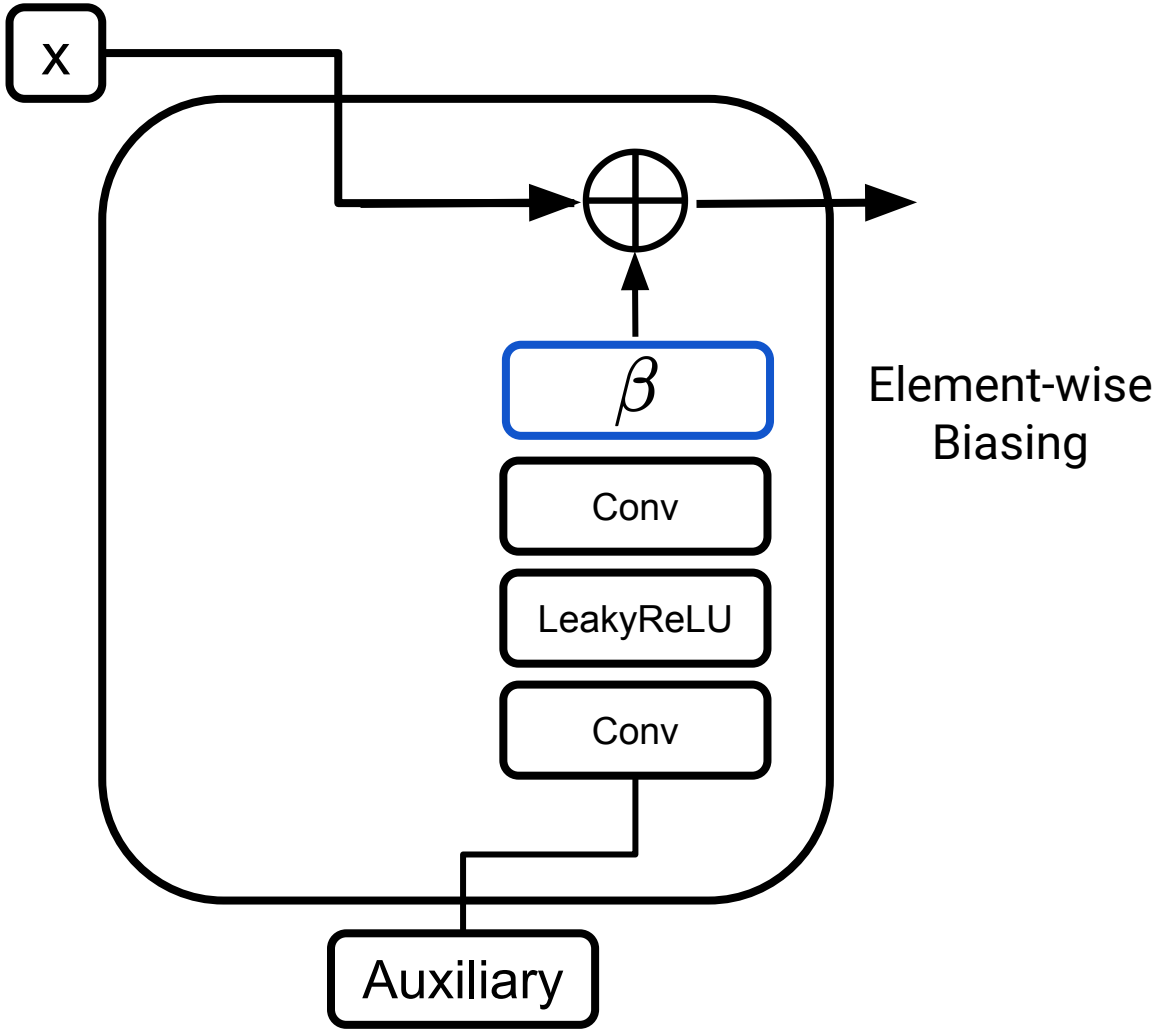

Lambda

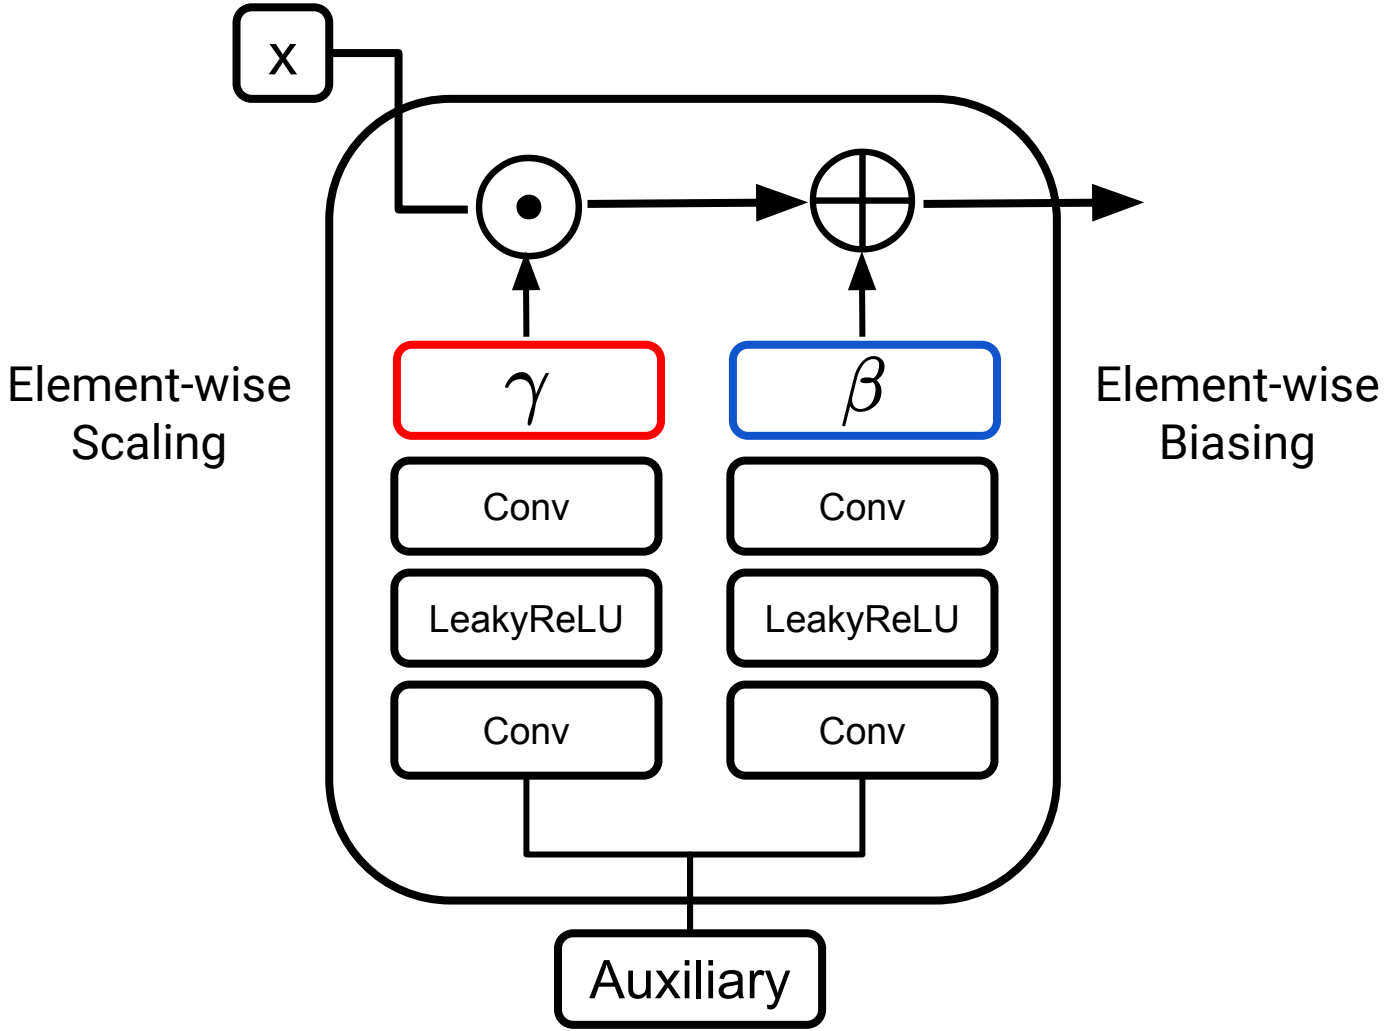
x
Element-wise
Scaling
γ
β
Element-wise
Biasing
Conv
LeakyReLU
Conv
Conv
LeakyReLU
Conv
Auxiliary

Lambda

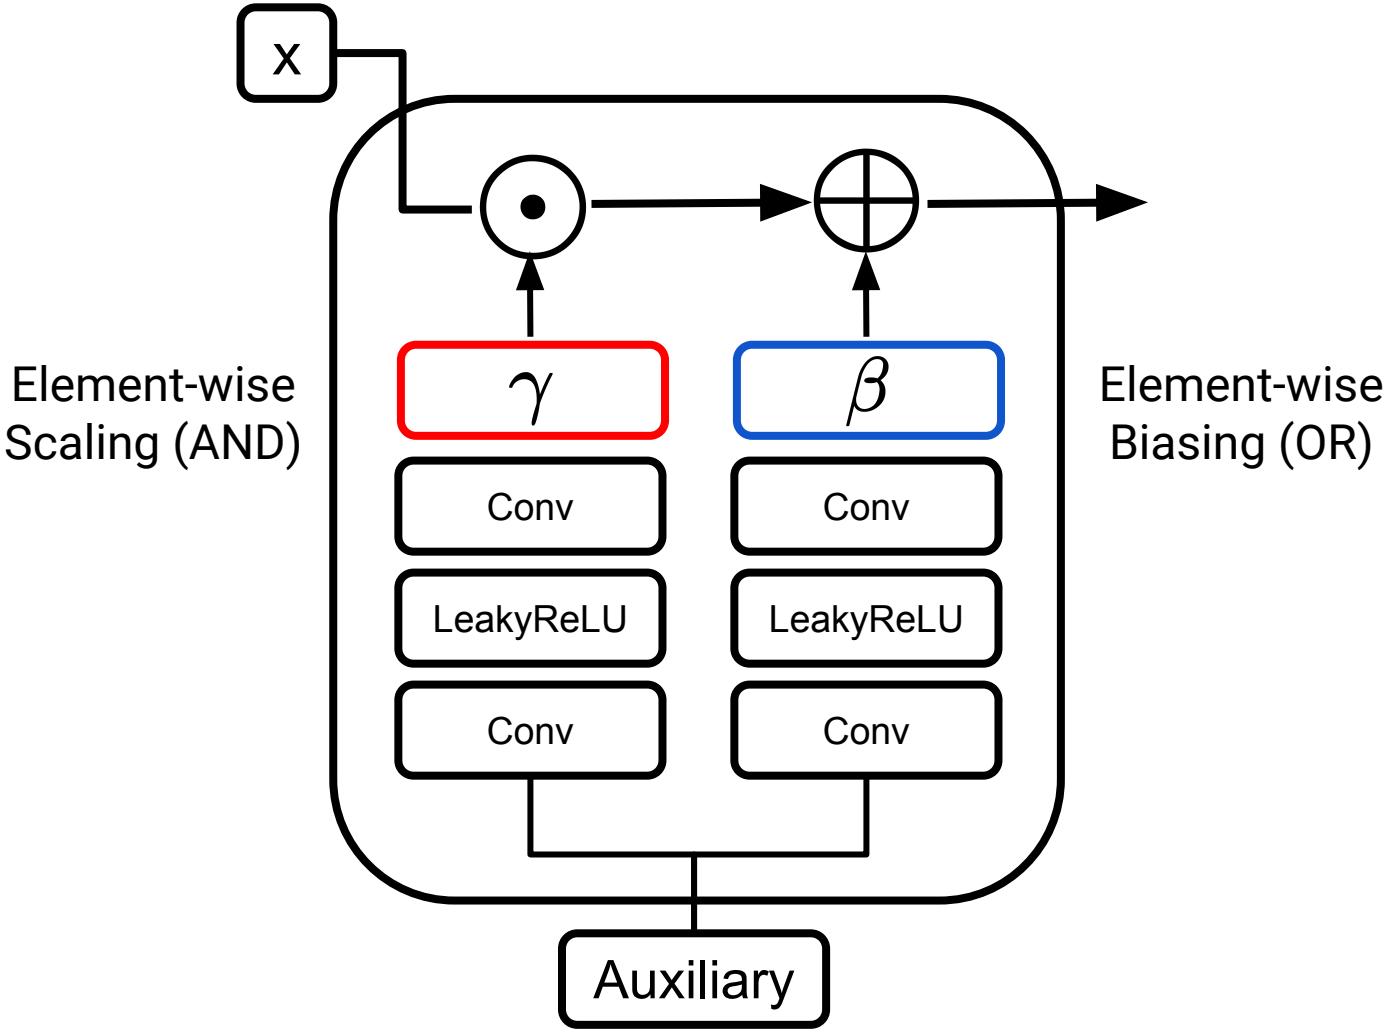
x
Element-wise
Scaling (AND)
γ
β
Element-wise
Biasing (OR)
Conv
LeakyReLU
Conv
Conv
LeakyReLU
Conv
Auxiliary

Lambda

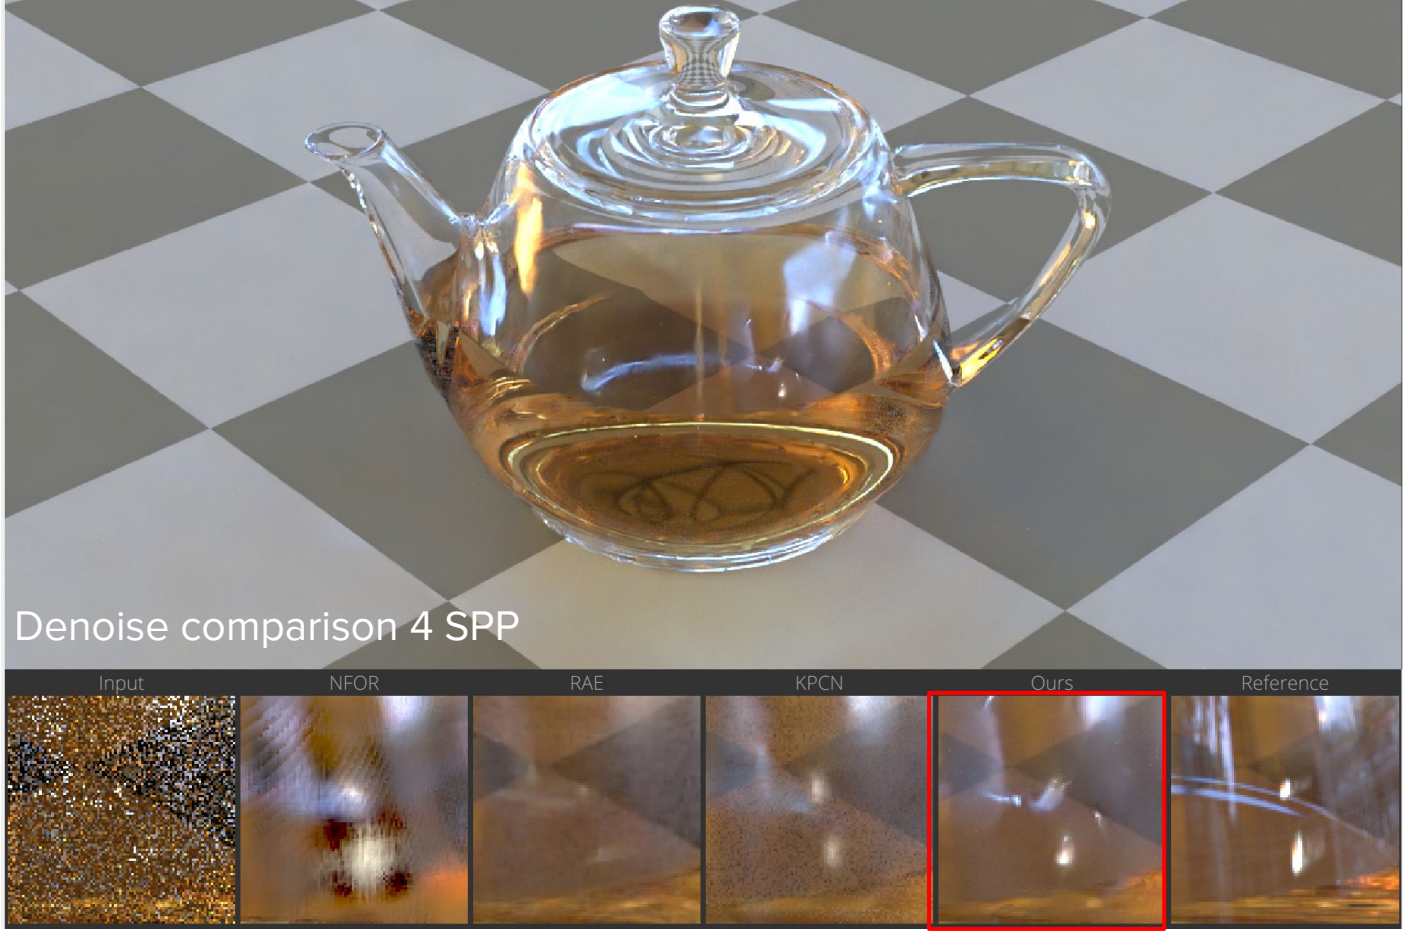

Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation

B Xu et al. Siggraph Asia 2019

Lambda

Value Networks

Denoising

4 SPP

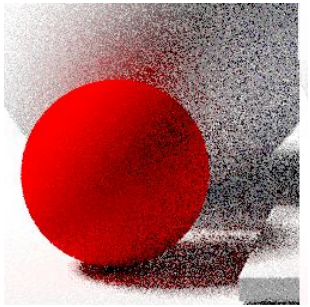

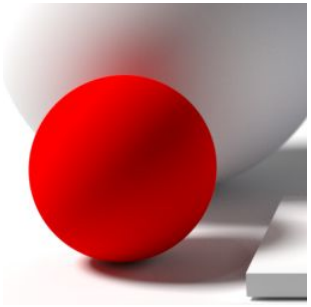

2^15 SPP

Policy Networks

Same SPP

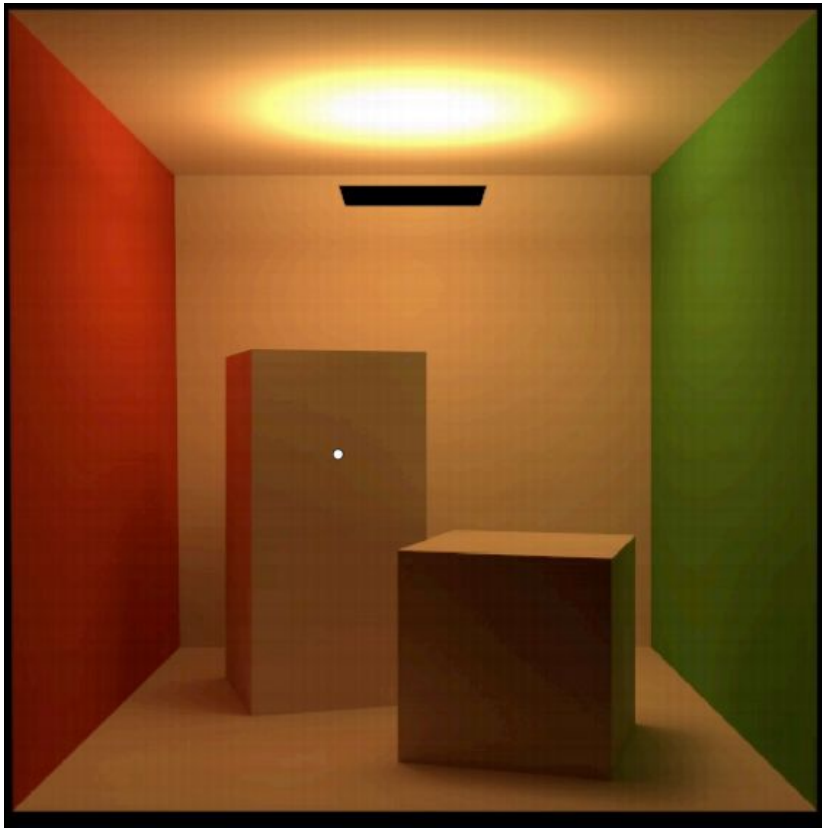

Neural Importance Sampling

Thomas Müller et al. ACM Transactions on Graphics 2019

Lambda

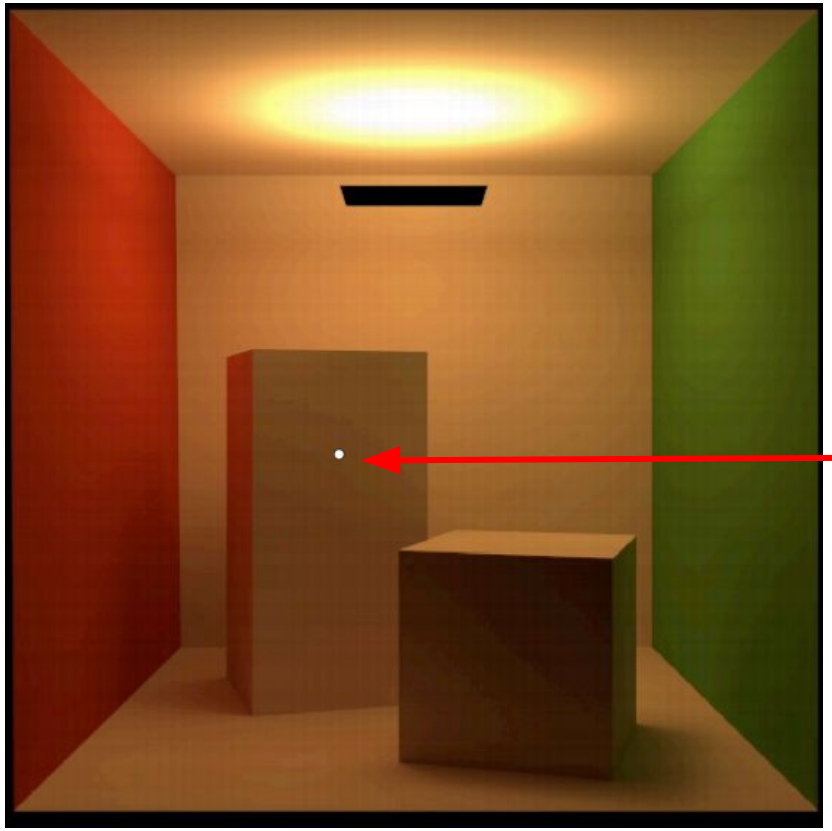

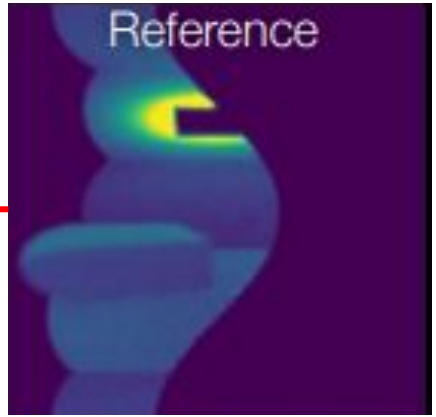

incidence radiance map

Neural Importance Sampling

Thomas Müller et al. ACM Transactions on Graphics 2019

Lambda

Z ⟶ $G(Z)$

$\mathcal{N}(0, \sigma^2)$

Lambda

$$Z \longrightarrow G(Z)$$

$$\mathcal{N}(0, \sigma^2)$$

$$[P, \omega_{in}, N]$$

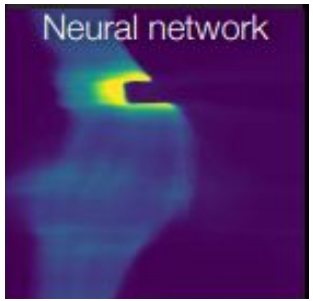

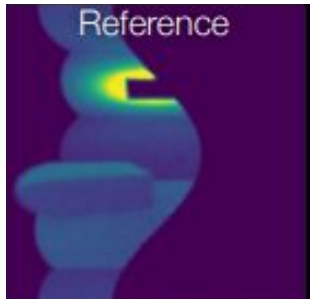

Neural Importance Sampling

Thomas Müller et al. ACM Transactions on Graphics 2019

Lambda

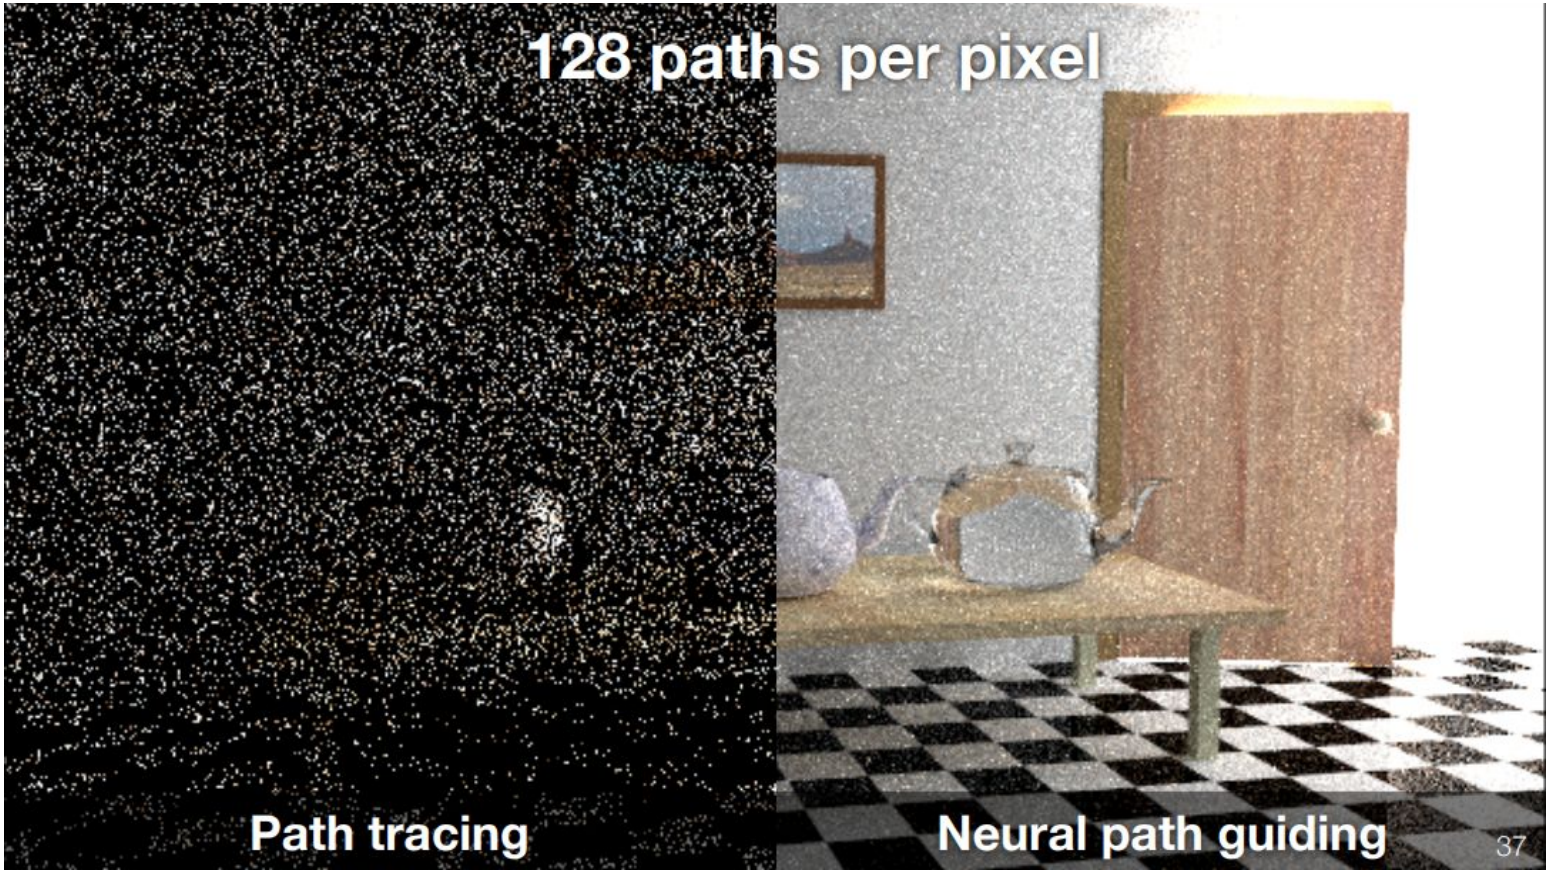

Neural Importance Sampling

Thomas Müller et al. ACM Transactions on Graphics 2019

Lambda

Sub-module

End-2-End

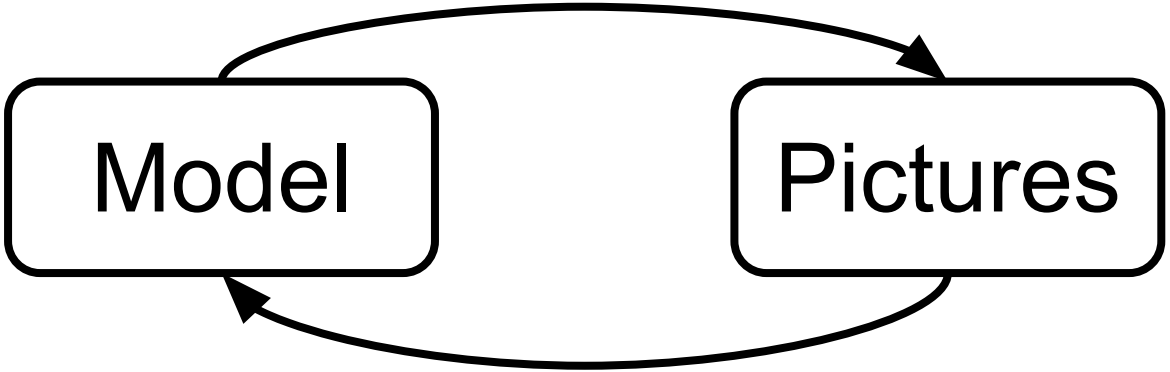

Differentiable Rendering

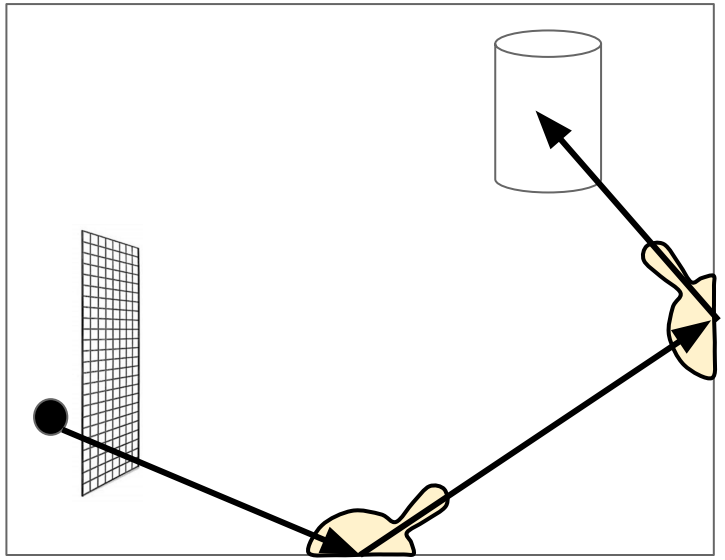

Ray Tracing
Image Centric

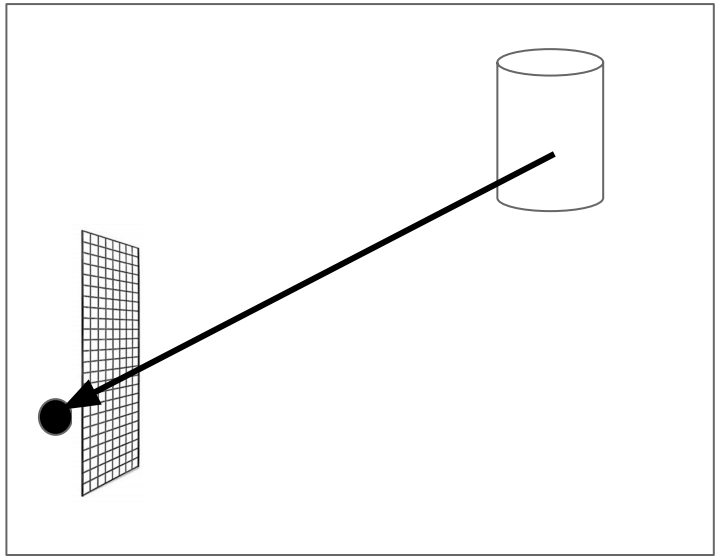

Rasterization
Object Centric

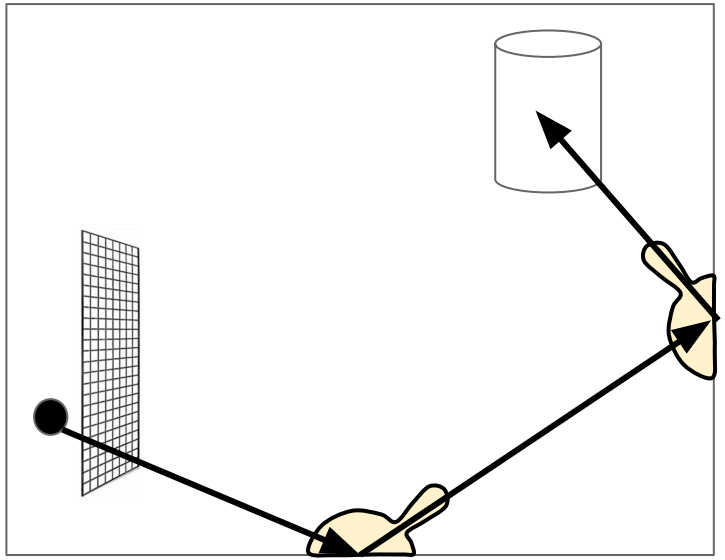

Ray Tracing
Image Centric

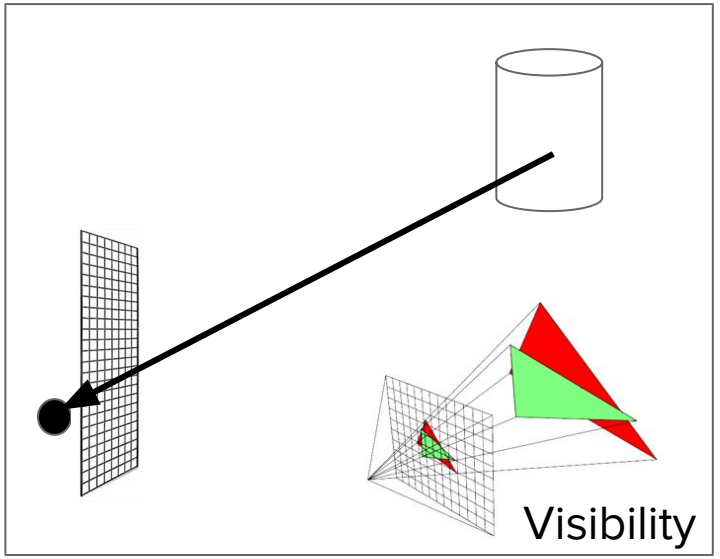

Rasterization
Object Centric

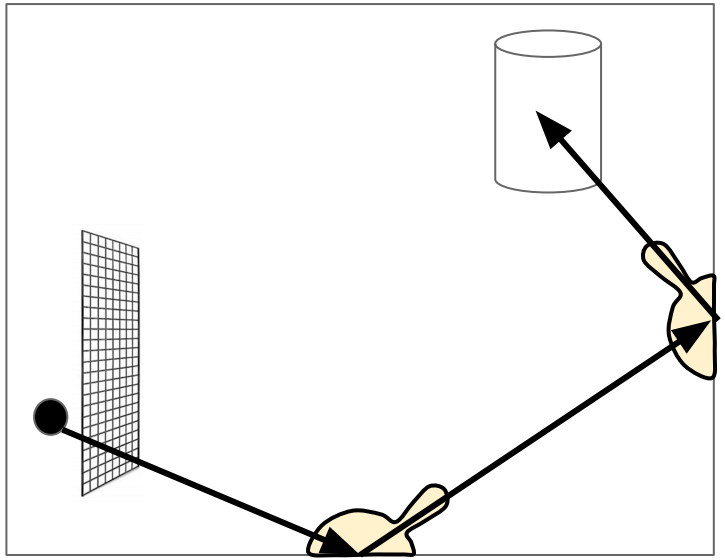

Ray Tracing
Image Centric

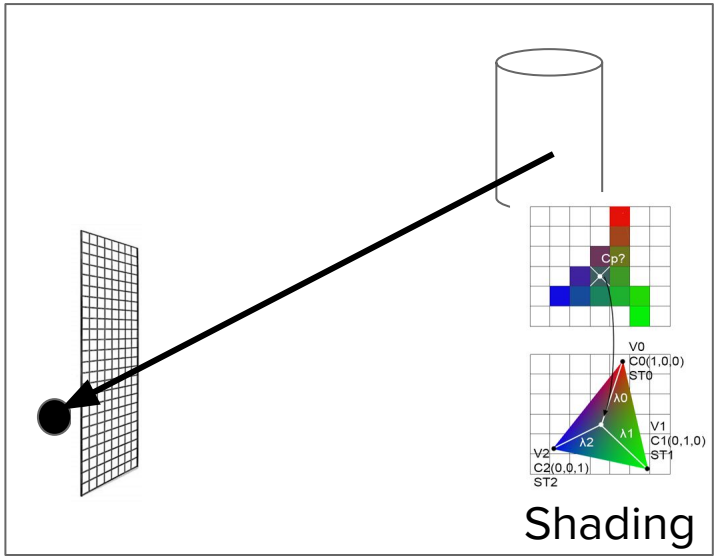

Rasterization
Object Centric

Lambda

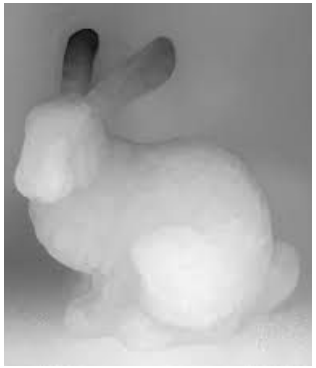

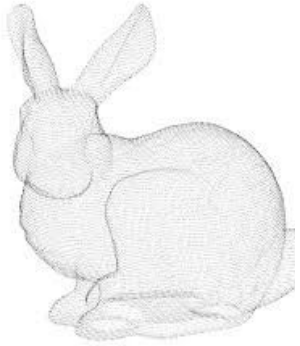

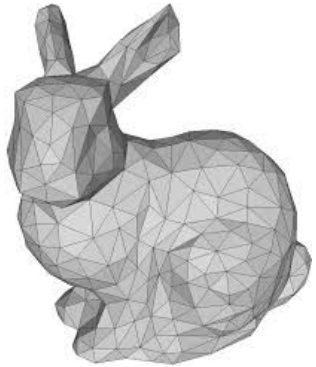

	Depth Map	Voxel	Point Cloud	Mesh
Memory	Good	Very Poor	Poor	Very Good
NN friendly	Great	Yes	No	Enemy

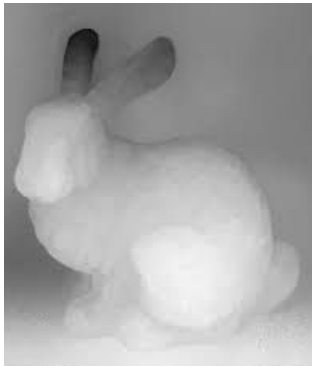

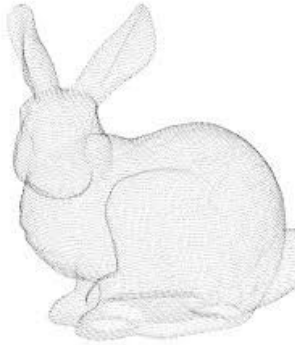

	Depth Map	Voxel	Point Cloud	Mesh
Memory	Good	Very Poor	Poor	Very Good
NN friendly	Great	Yes	No	Enemy

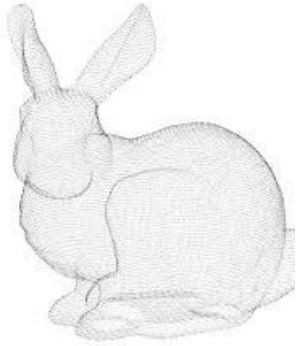

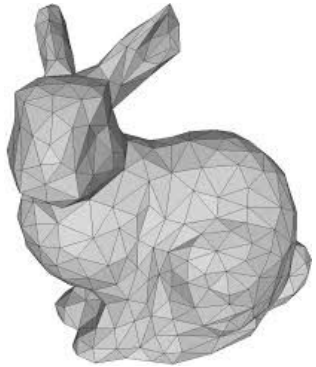

	Depth Map	Voxel	Point Cloud	Mesh
Memory	Good	Very Poor	Poor	Very Good
NN friendly	Great	Yes	No	Enemy

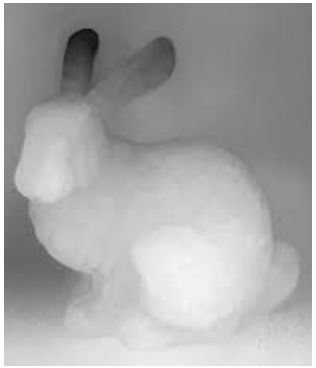

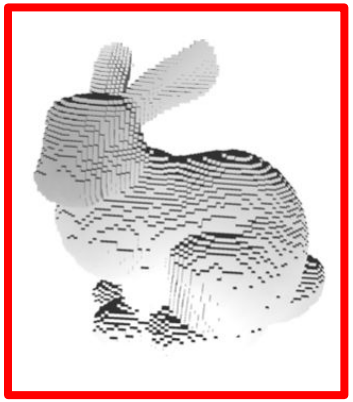

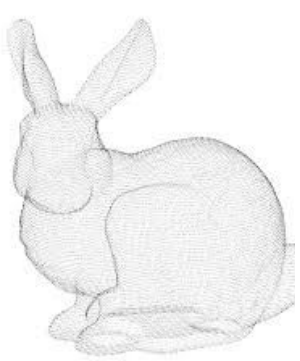

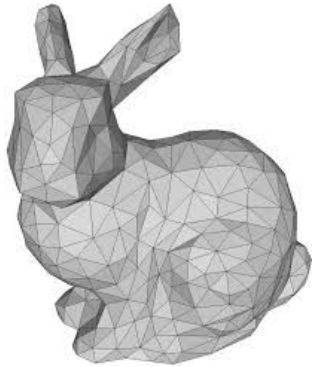

	Depth Map	Voxel	Point Cloud	Mesh
Memory	Good	Very Poor	Poor	Very Good
NN friendly	Great	Yes	No	Enemy

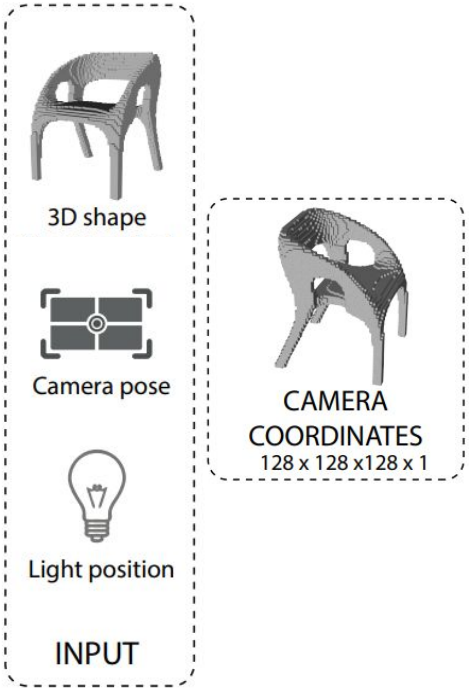

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

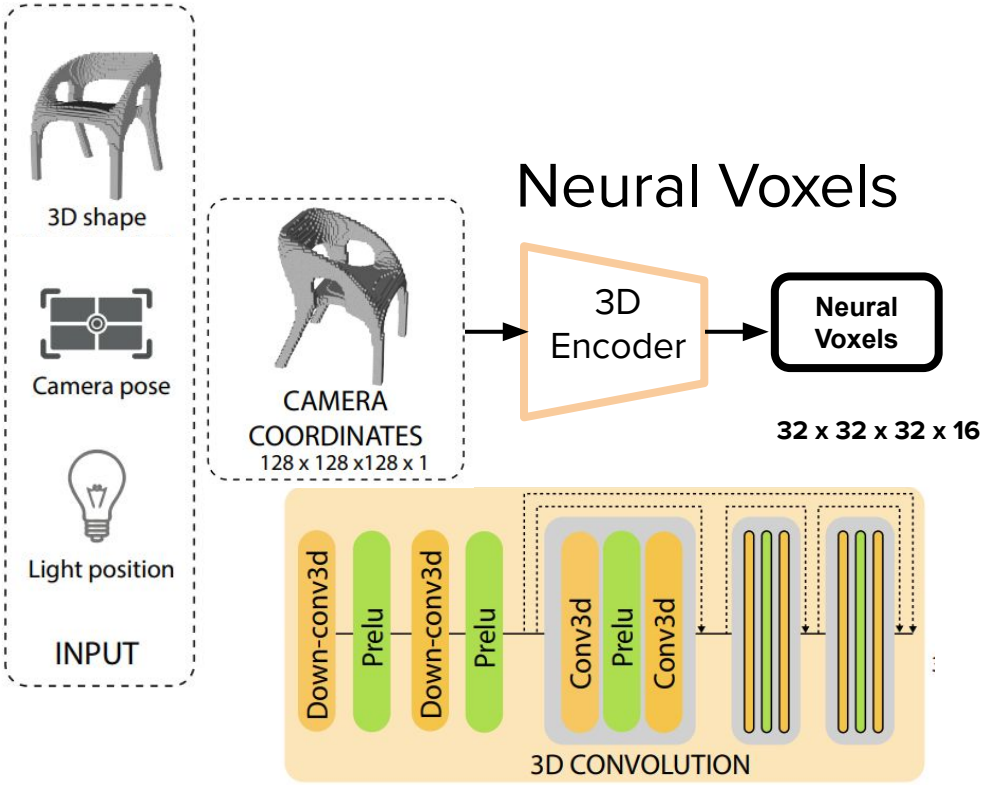

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

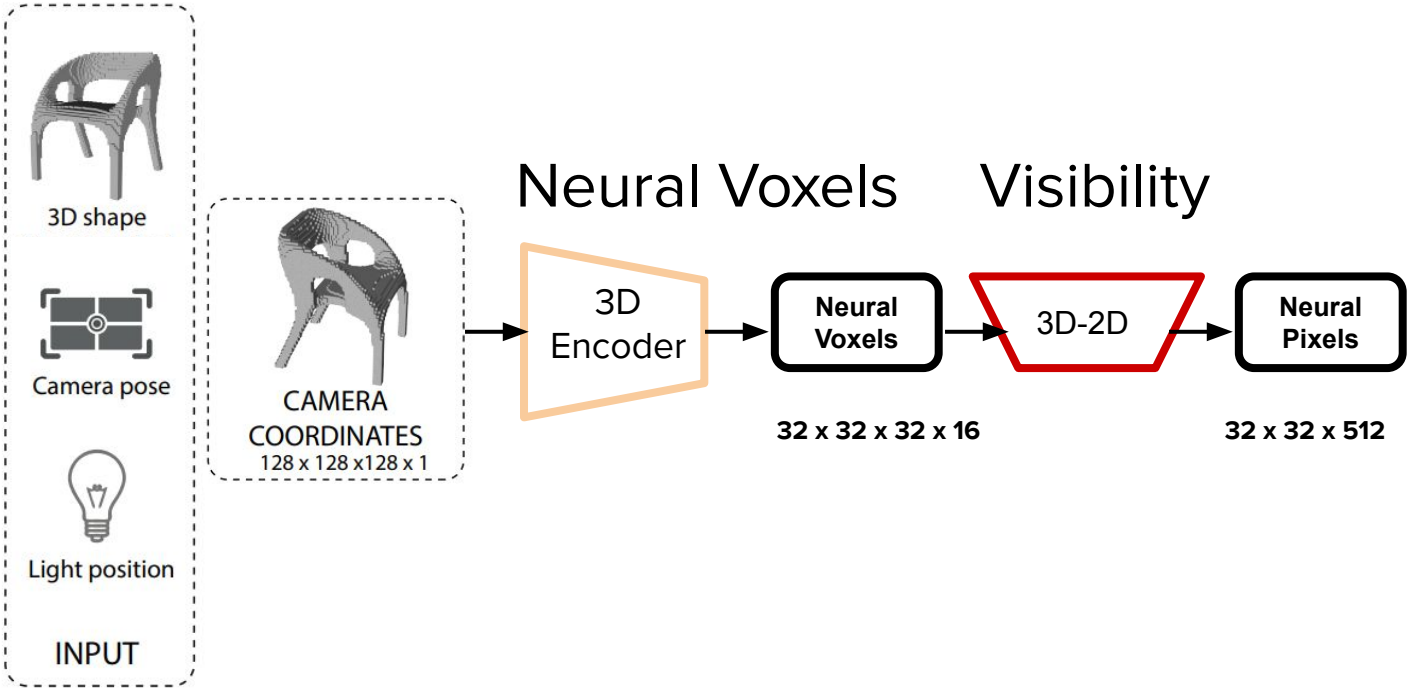

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

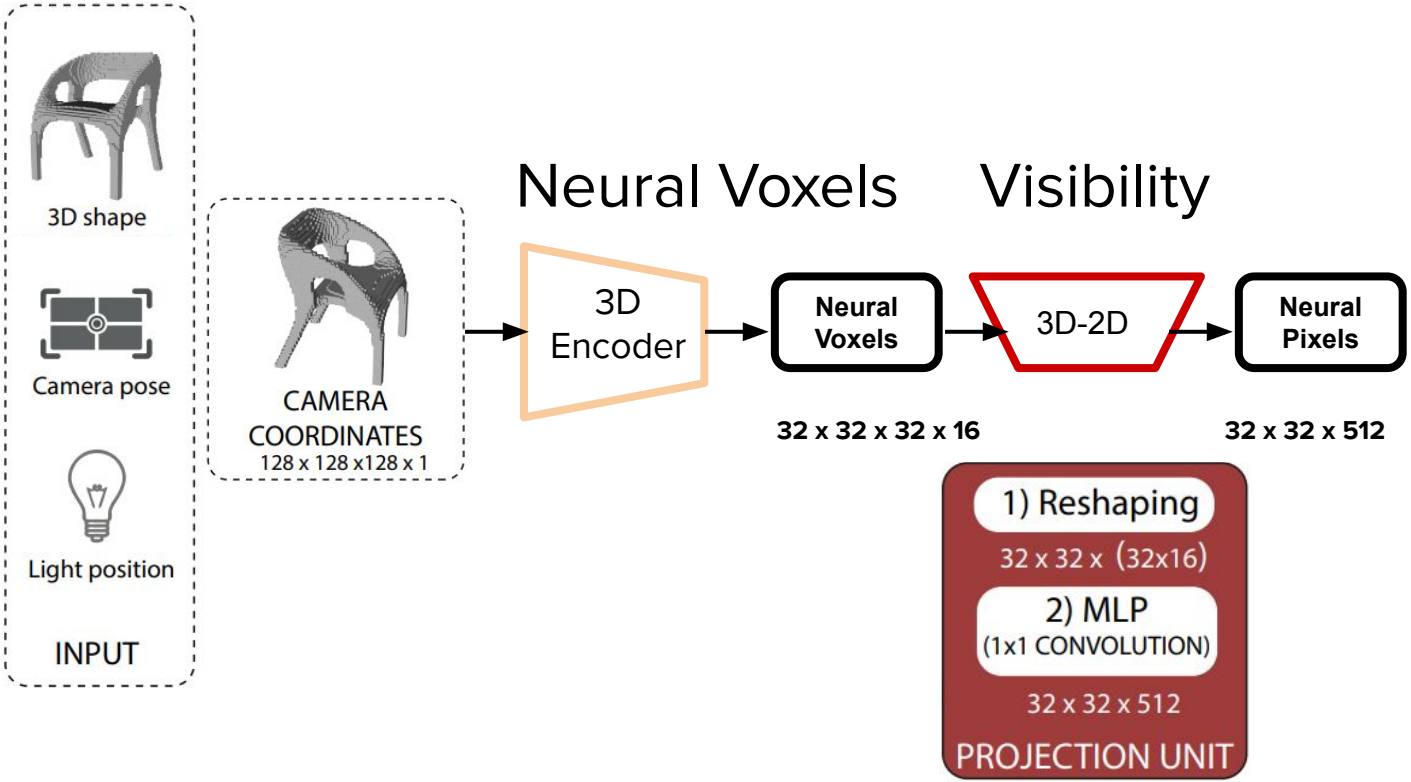

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

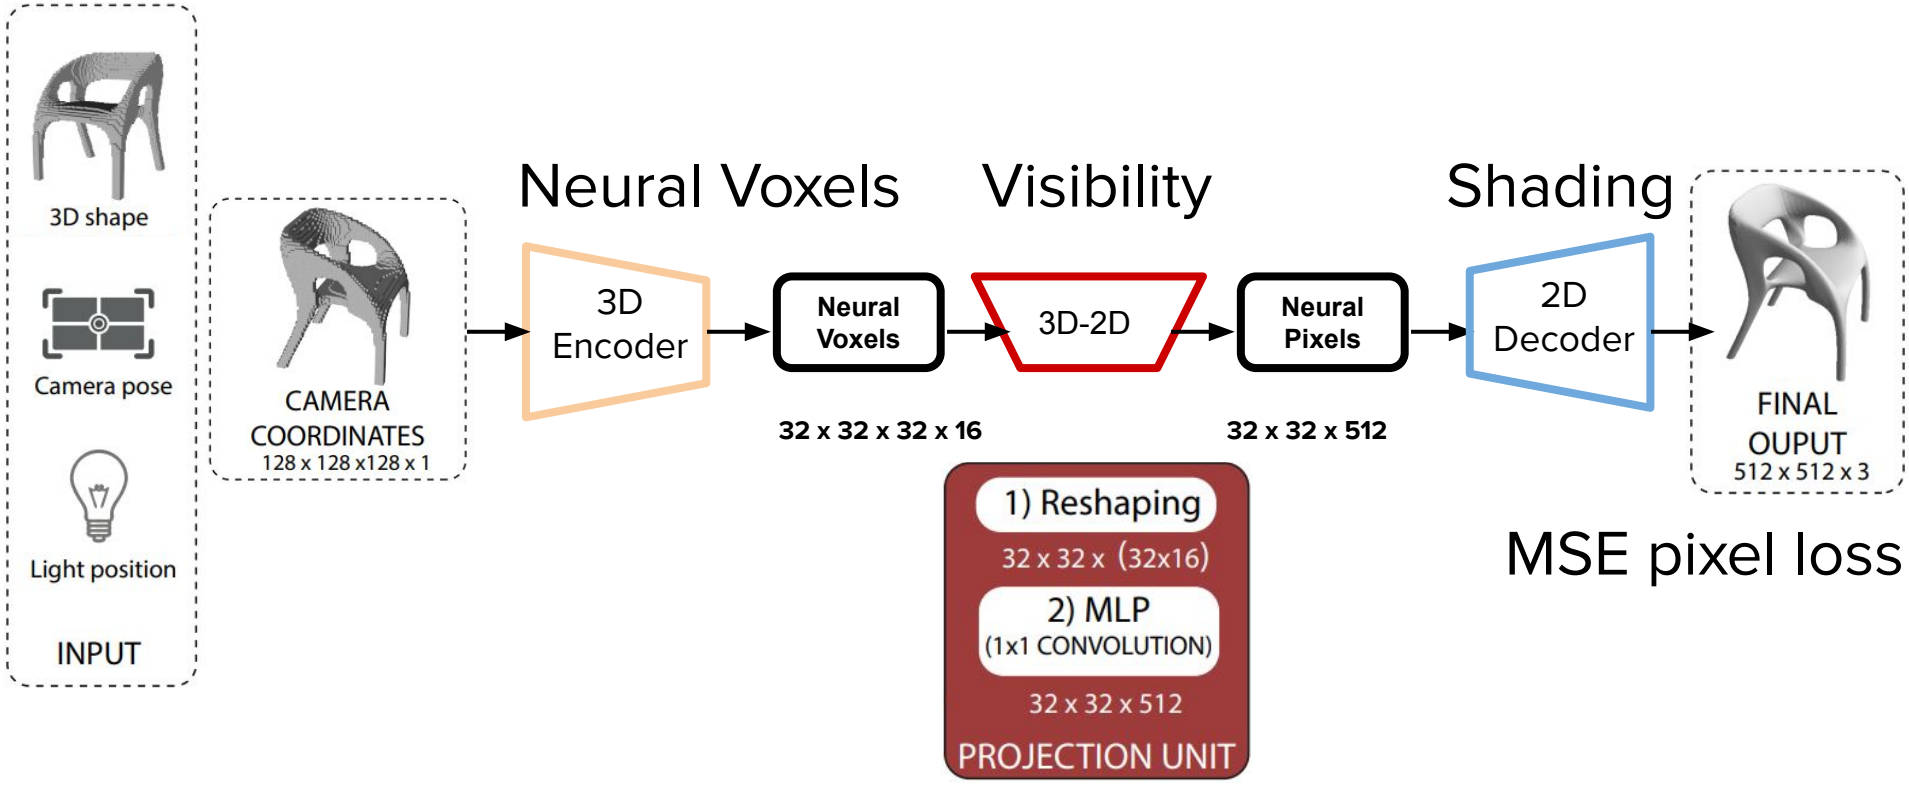

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

Contour

Toon

Ambient Occlusion

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

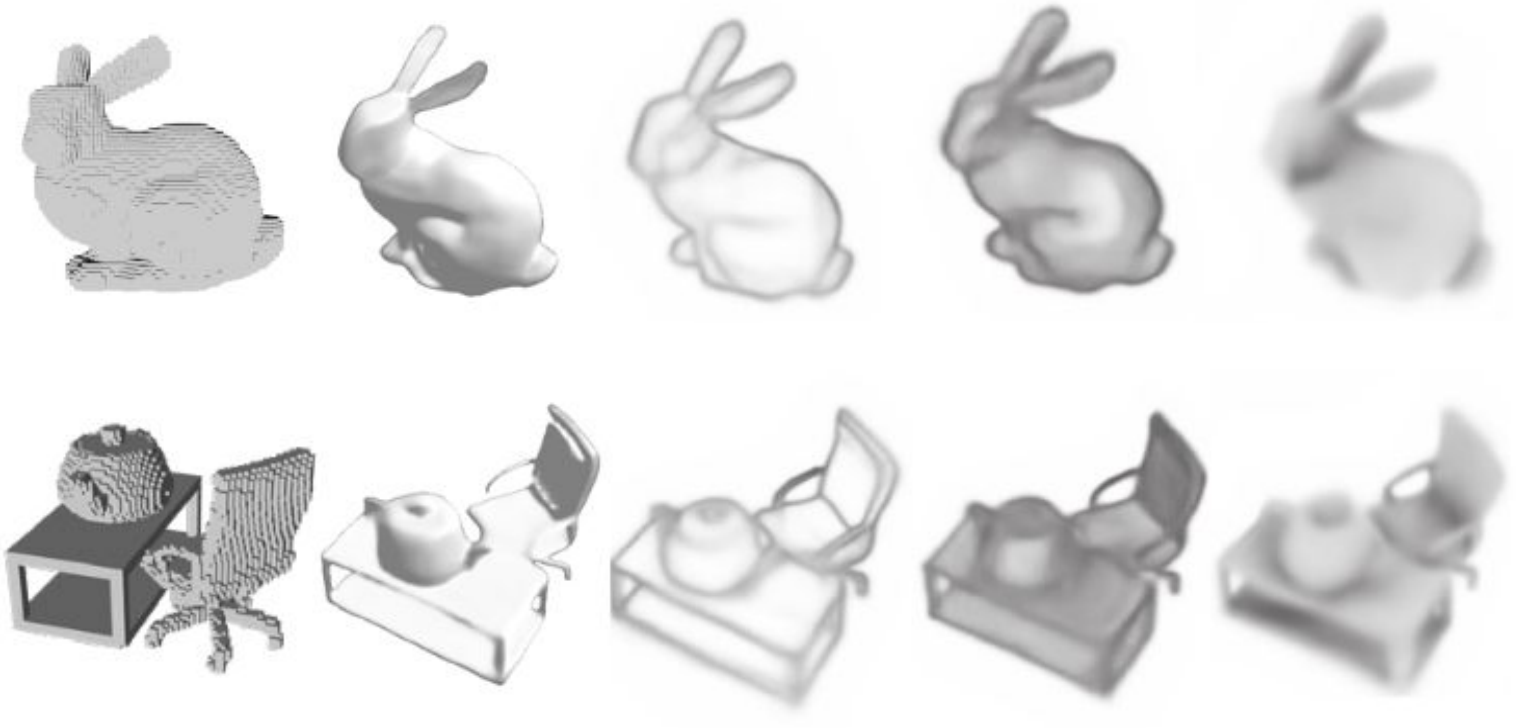

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

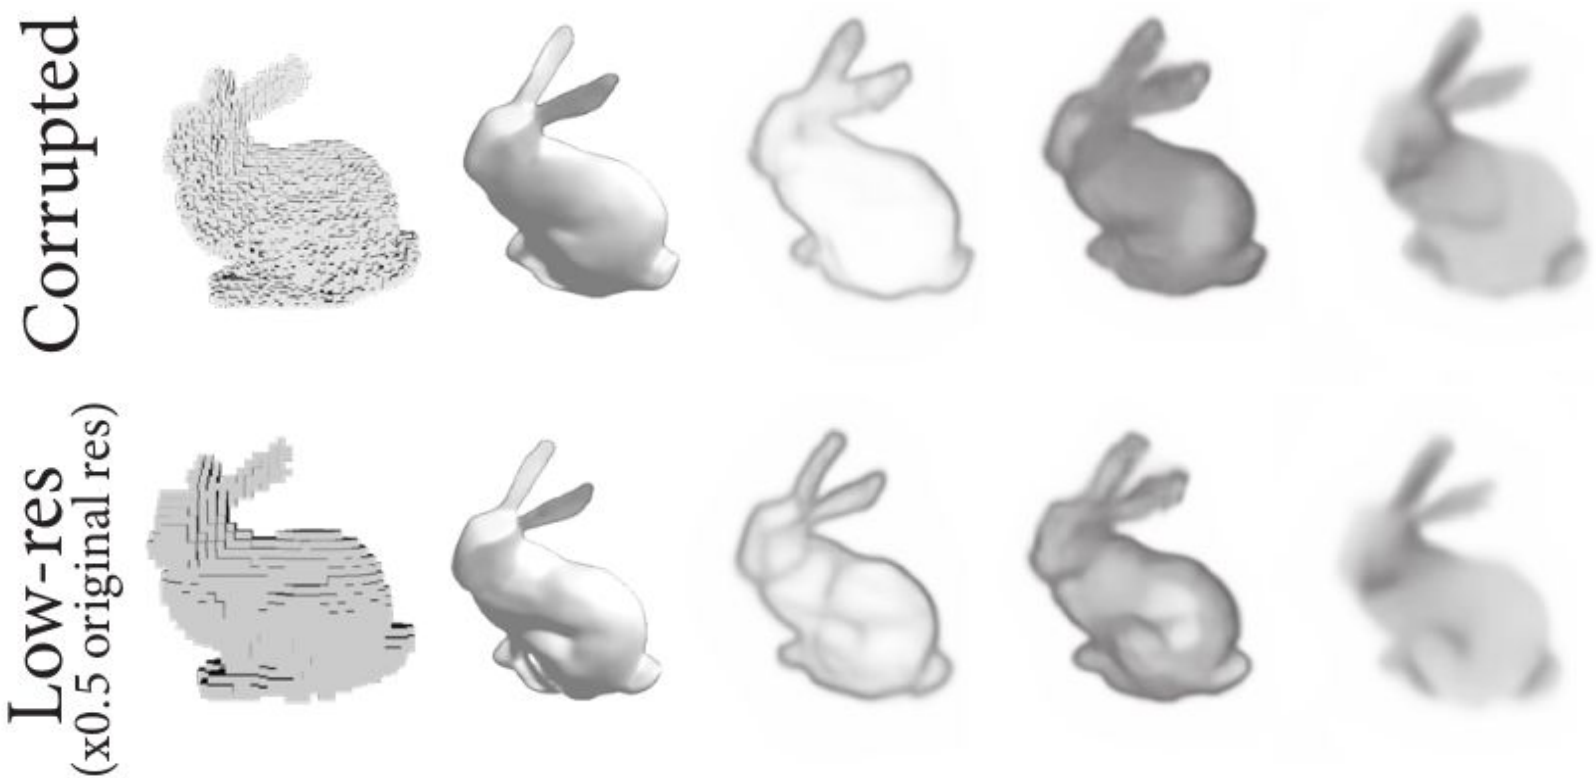

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

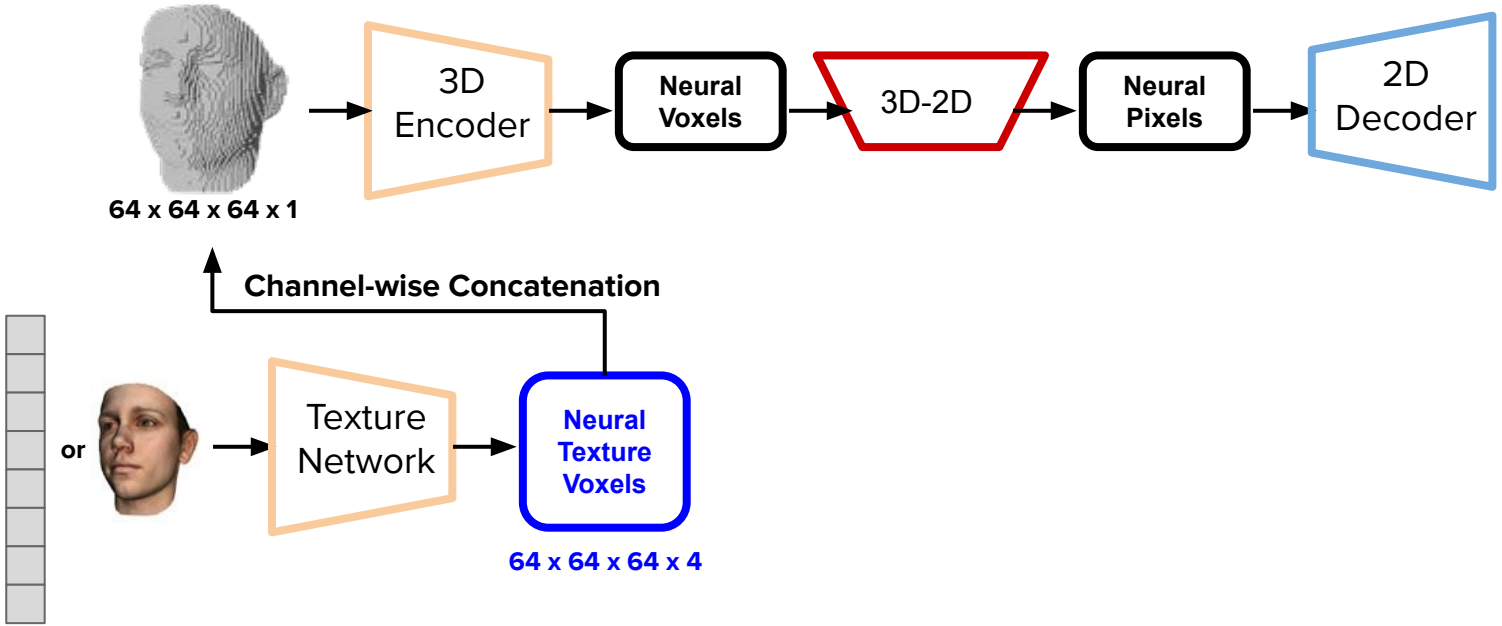

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

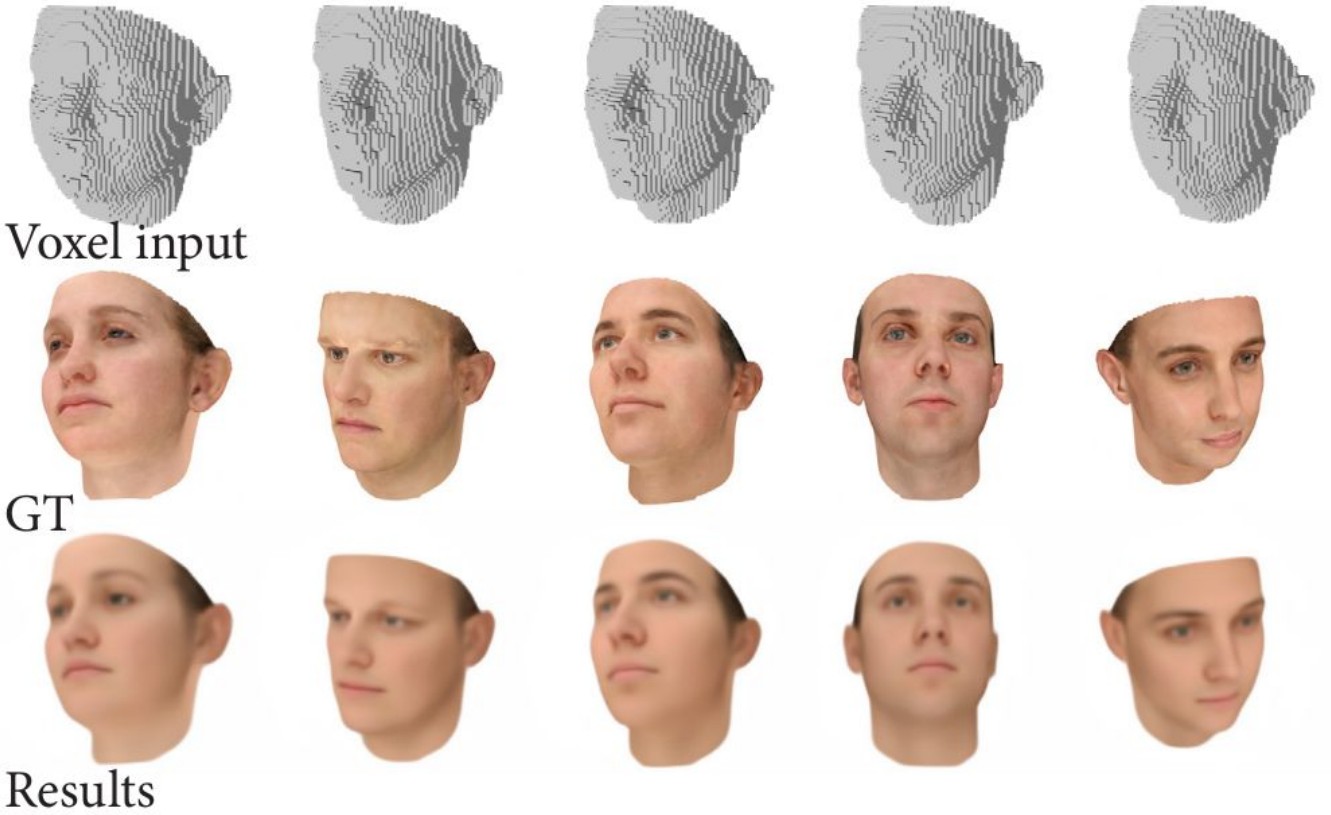

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

Same shape, different textures

Same texture, different shapes

RenderNet: A deep convolutional network for differentiable rendering from 3D shapes

Thu Nguyen-Phuoc et al. NeurIPS 2018

Lambda

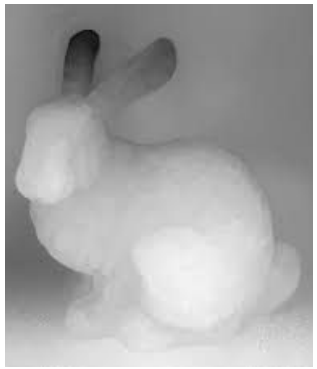

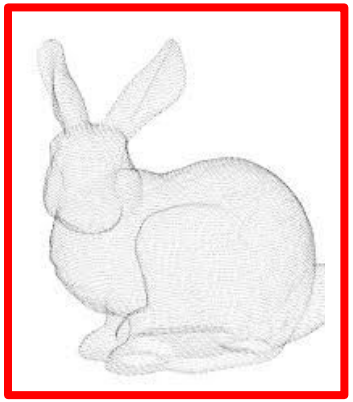

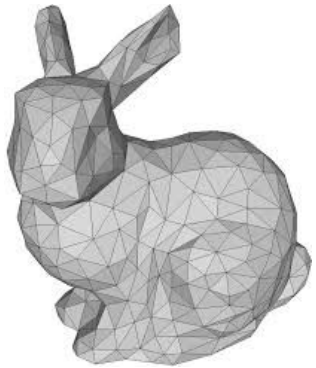

	Depth Map	Voxel	Point Cloud	Mesh
Memory	Good	Very Poor	Poor	Very Good
NN friendly	Great	Yes	No	Enemy

Rasterization a RGB point cloud

Neural Point-Based Graphics

KA Aliev et al, arxiv 2019

Lambda

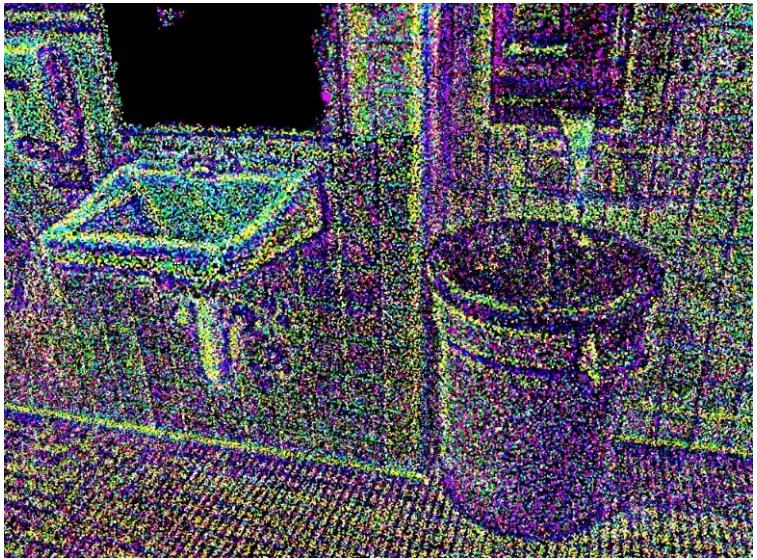

Rasterization a neural point cloud

(First three PCA dimensions of the neural descriptor)

Neural Point-Based Graphics

KA Aliev et al, arxiv 2019

Lambda

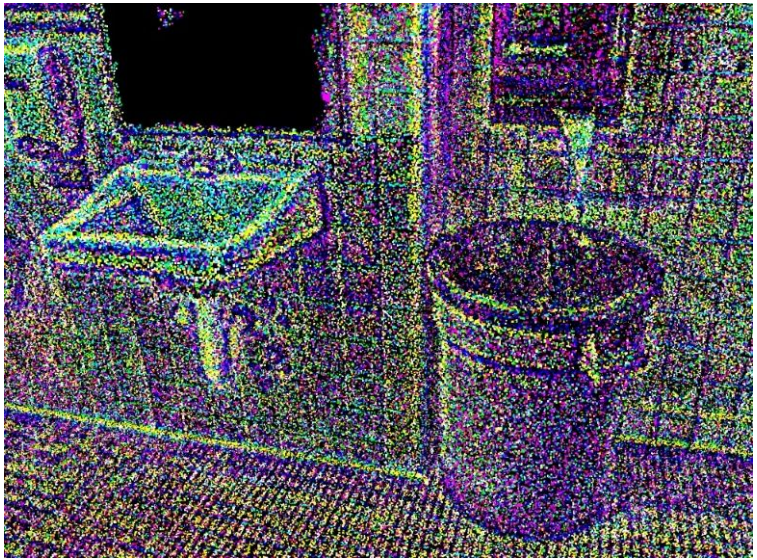

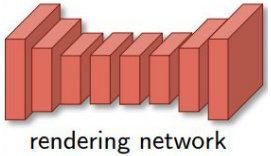

Rasterization a neural point cloud

(First three PCA dimensions of the neural descriptor)

Neural Point-Based Graphics

KA Aliev et al, arxiv 2019

Lambda

RBG rasterization

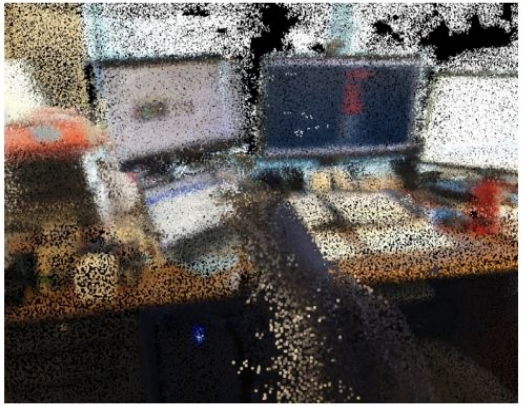

Neural rasterization

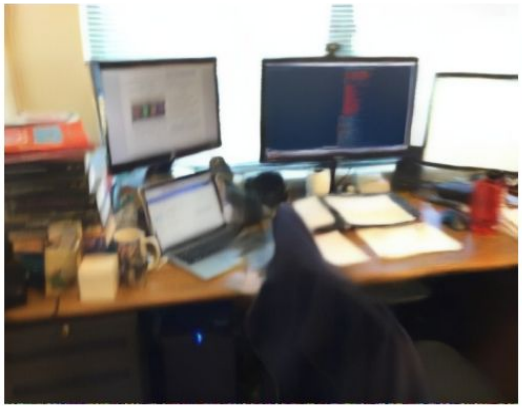
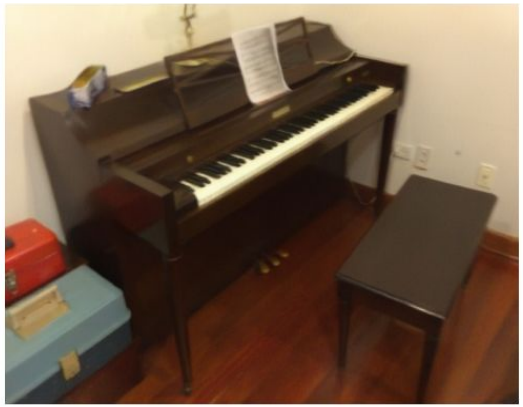
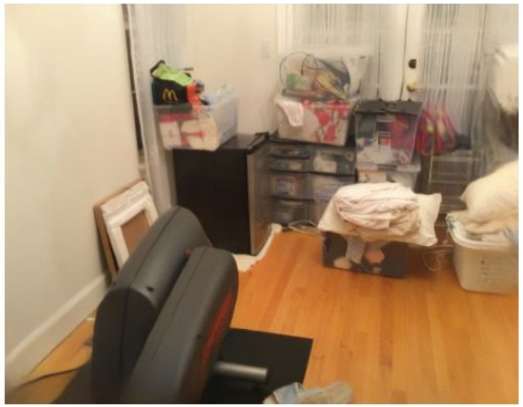

Neural Point-Based Graphics

KA Aliev et al, arxiv 2019

Lambda

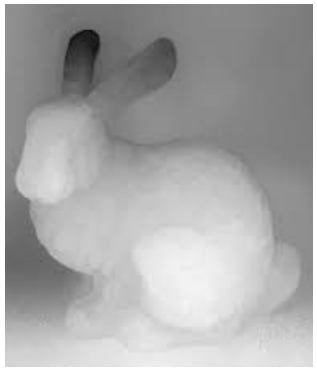

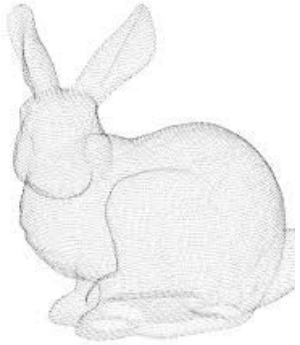

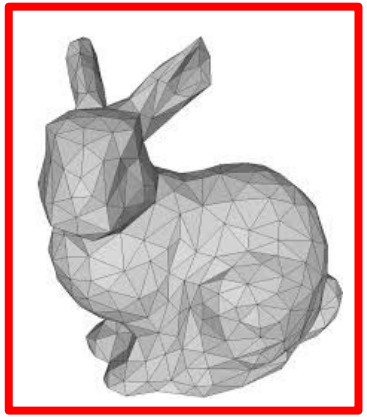

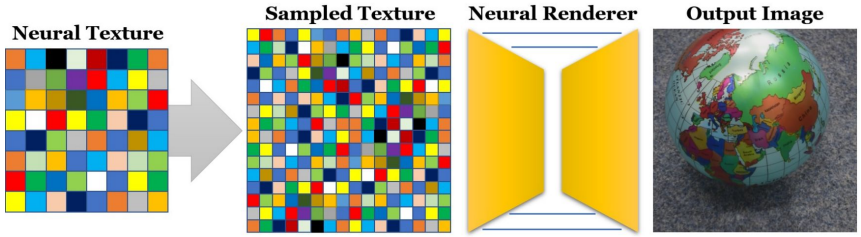

Deferred Neural Rendering:
Image Synthesis using Neural Textures

J Thies et al, Siggraph 2019

Neural 3D Mesh Renderer

H Kato et al, CVPR 2018

Lambda

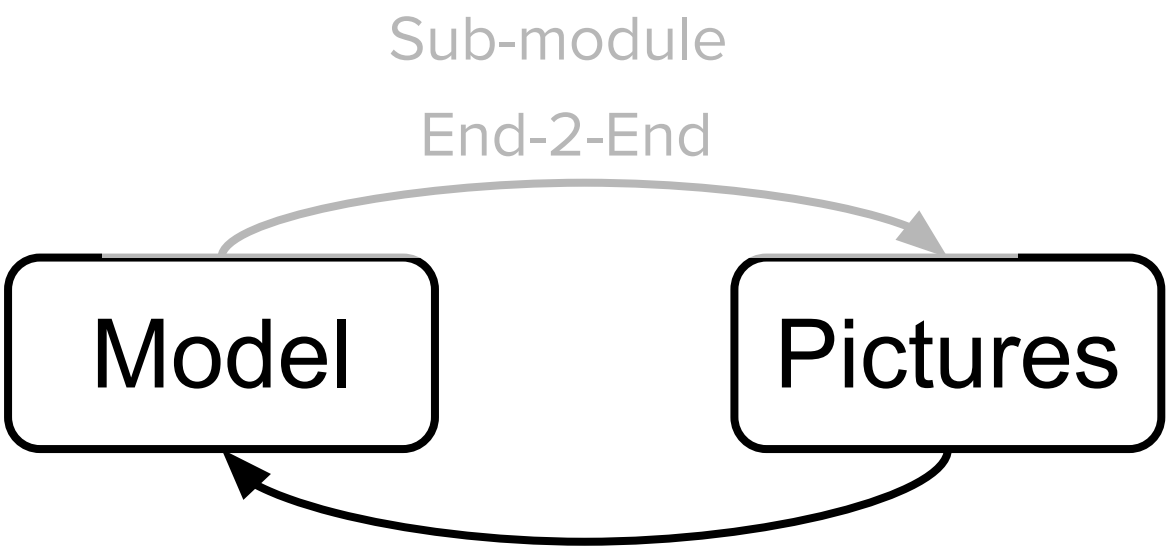
Sub-module
End-2-End
Model
Pictures

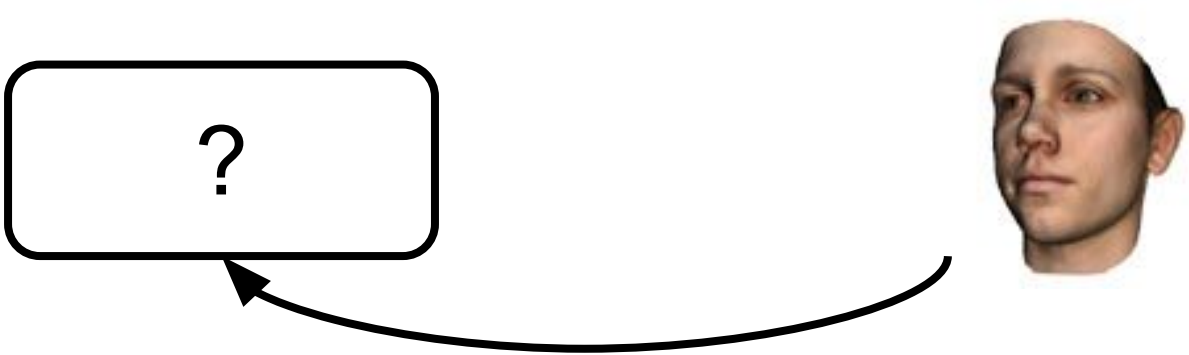

Lambda

Approximation

Target

Lambda

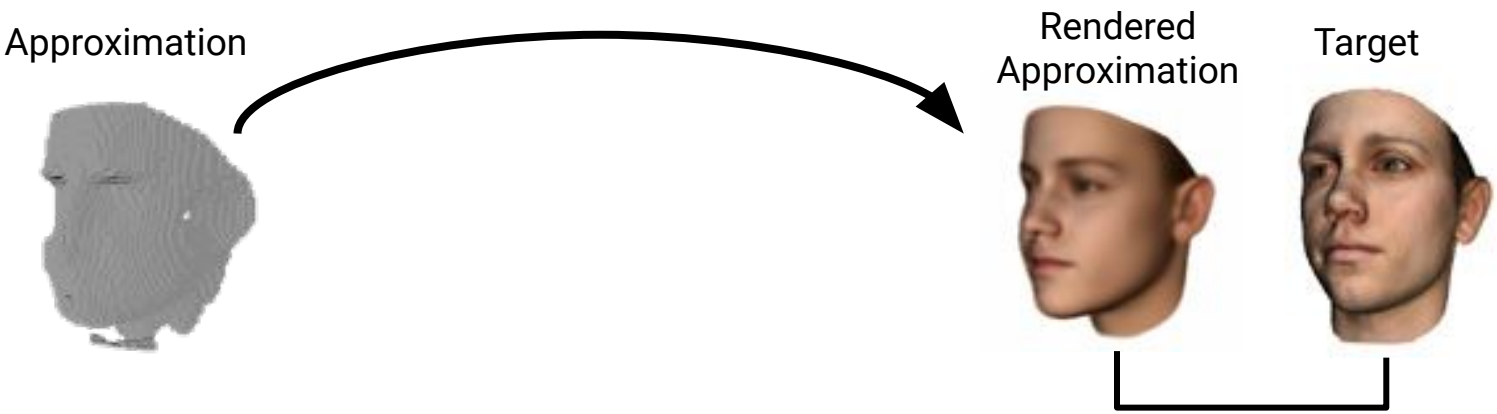

Lambda

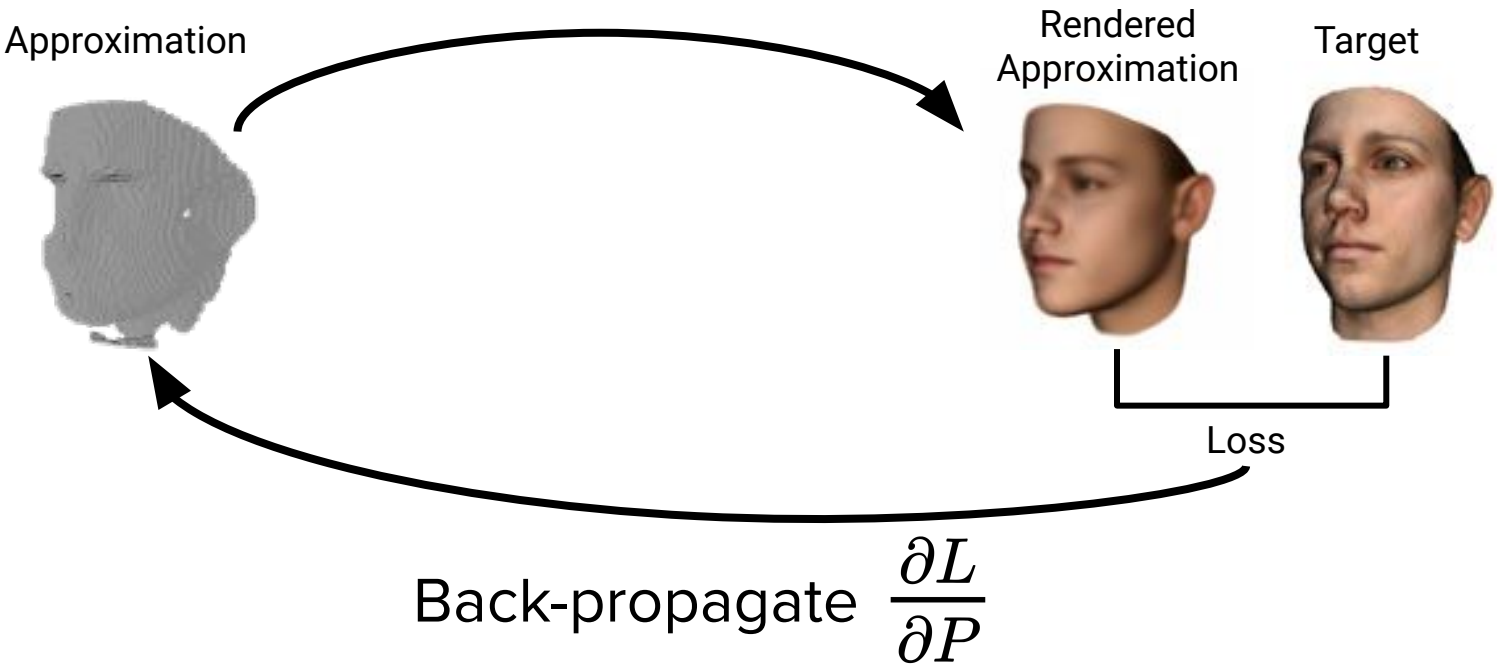
Approximation
Rendered Approximation
Target
Loss
Back-propagate $\frac{\partial L}{\partial P}$

Lambda

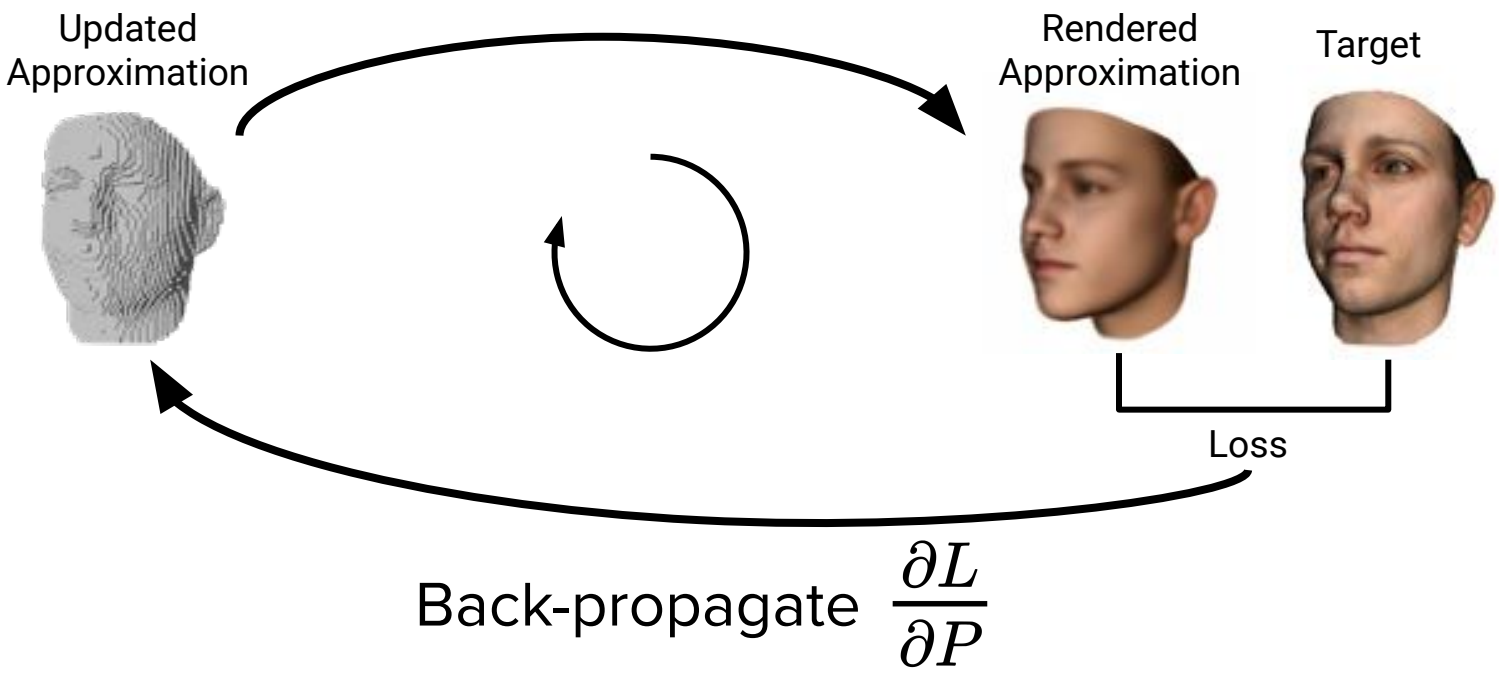
Updated
Approximation
Rendered
Approximation
Target
Loss
Back-propagate $\frac{\partial L}{\partial P}$

Lambda

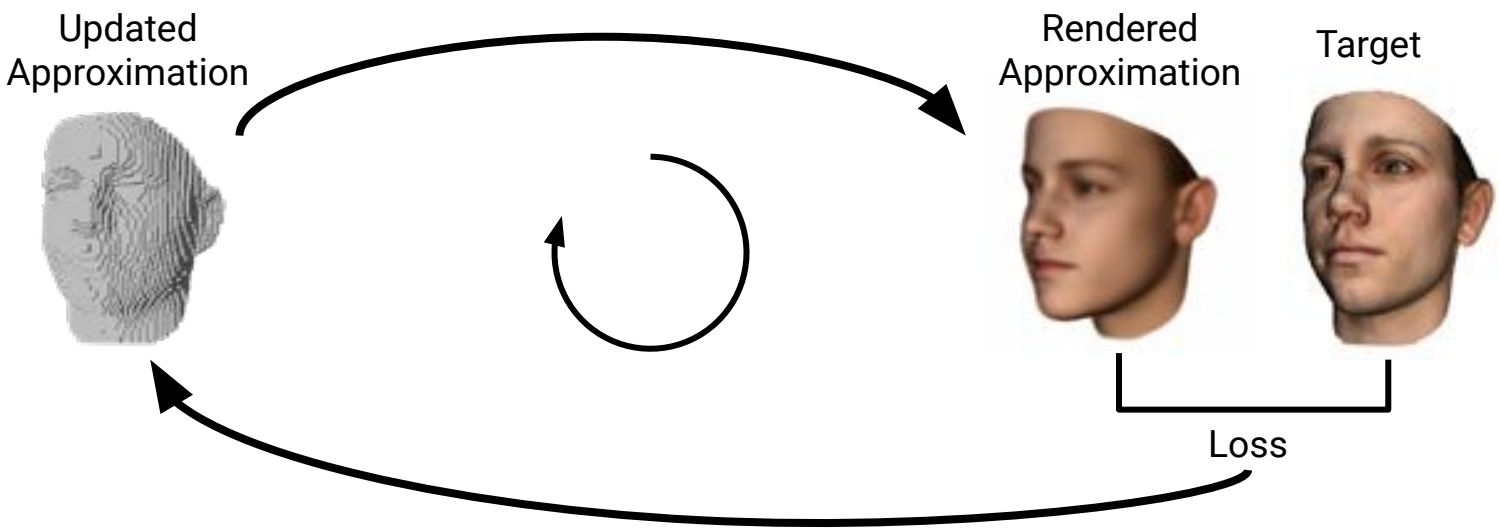

Back-propagate $\frac{\partial L}{\partial P}$

For Free

Lambda

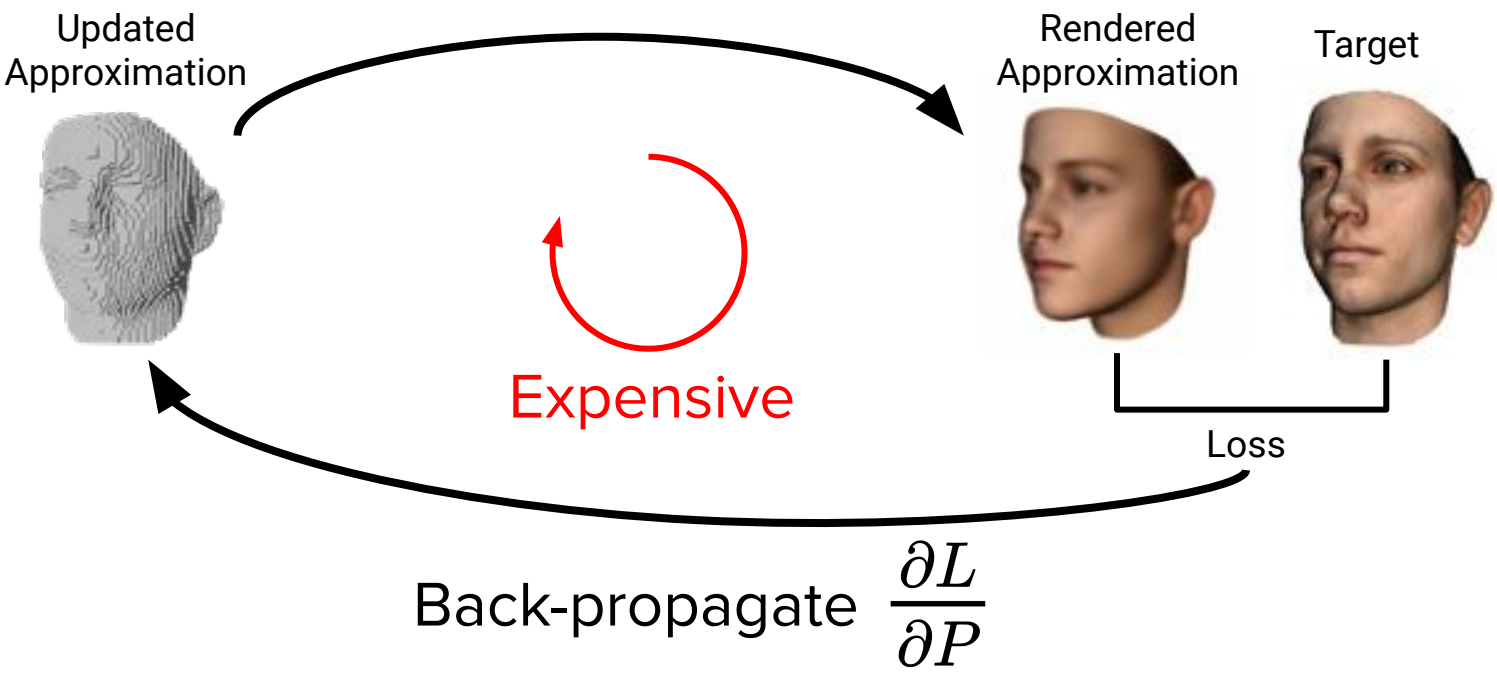
Updated
Approximation
Rendered
Approximation
Target
Expensive
Loss
Back-propagate $\frac{\partial L}{\partial P}$

Lambda

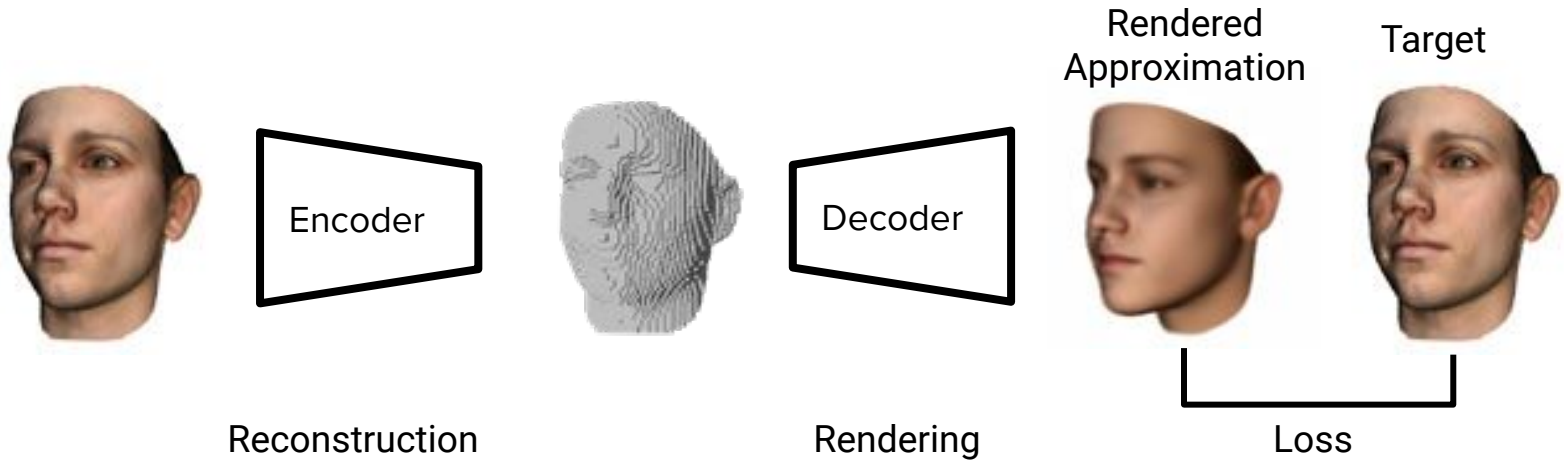
Rendered
Approximation
Target
Encoder
Decoder
Reconstruction
Rendering
Loss

Lambda

Inductive Bias: Separate Appearance from Pose

Human perception imposes coordinate frame on objects

Lambda

Learning 3D representation from natural images without 3D supervision

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

Conditional GANs

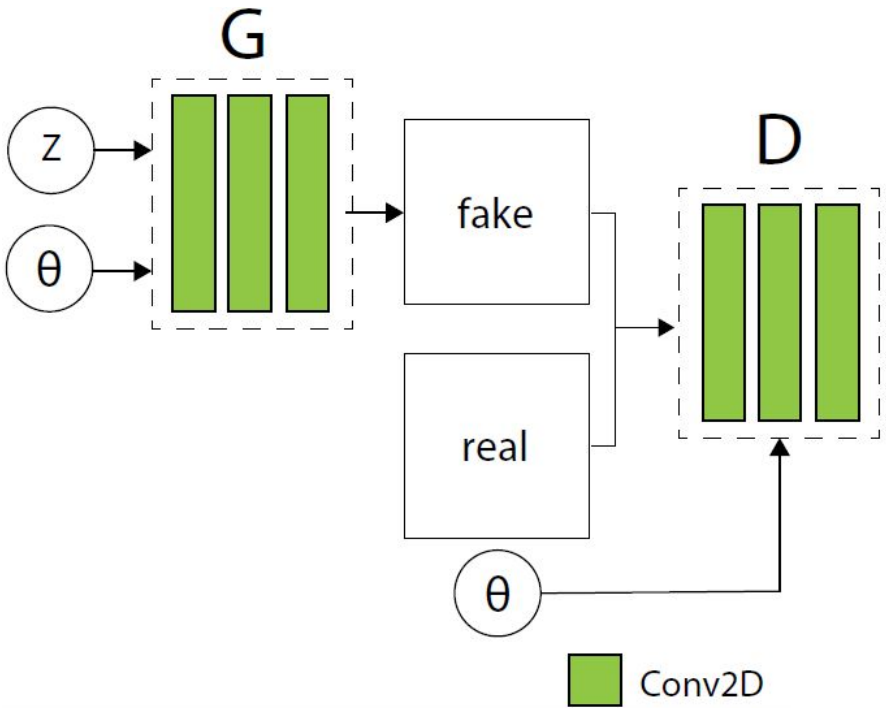

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Conditional GANs

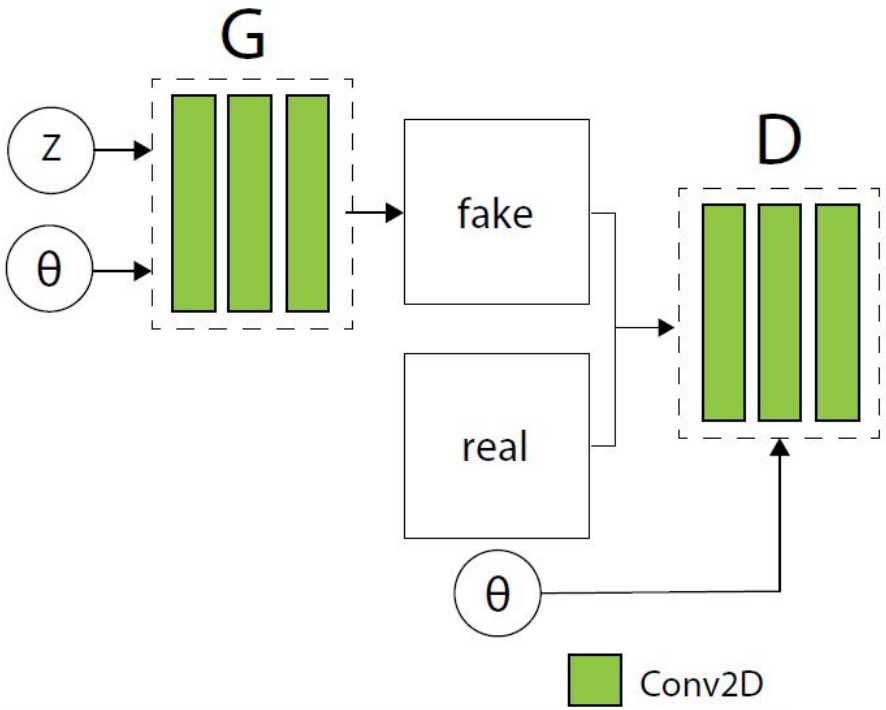

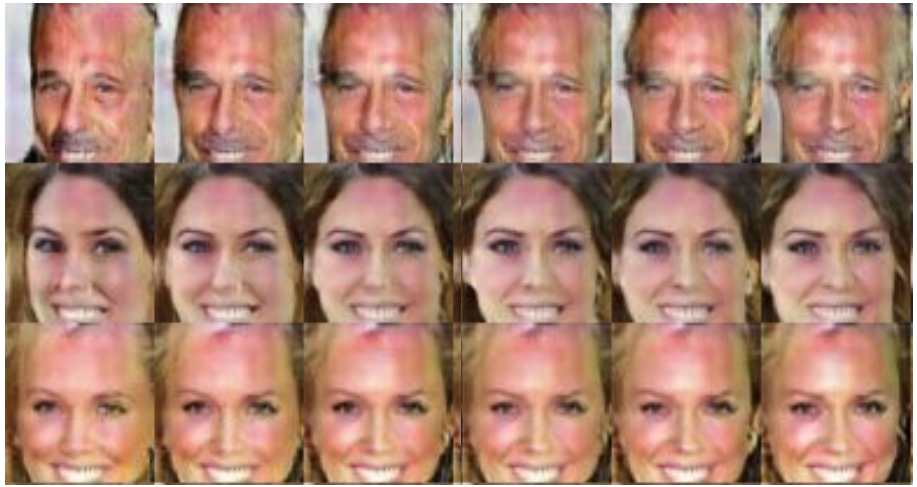

Info GANs

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

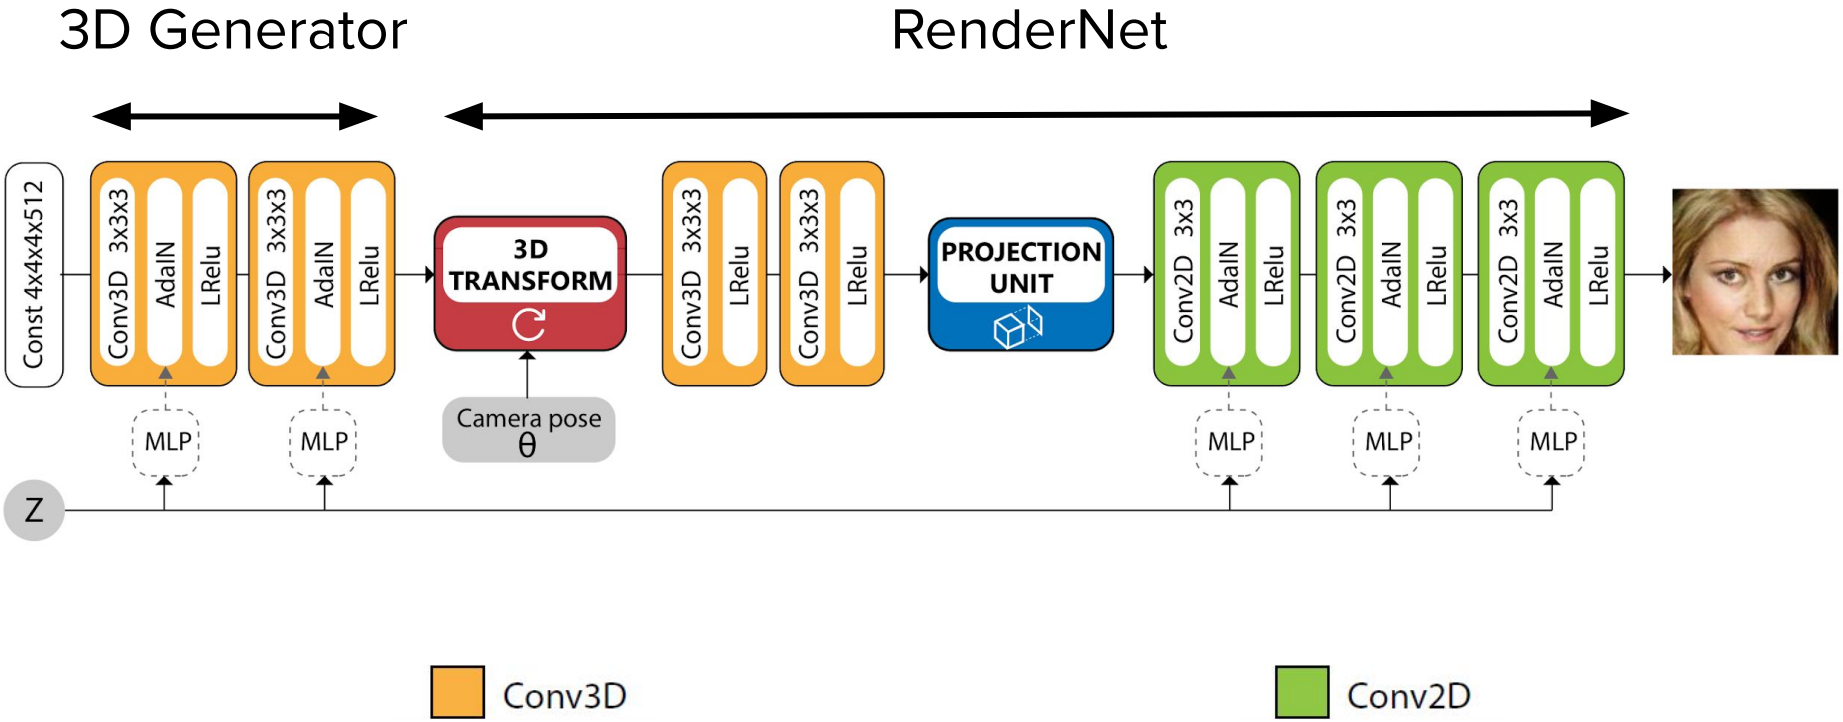

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

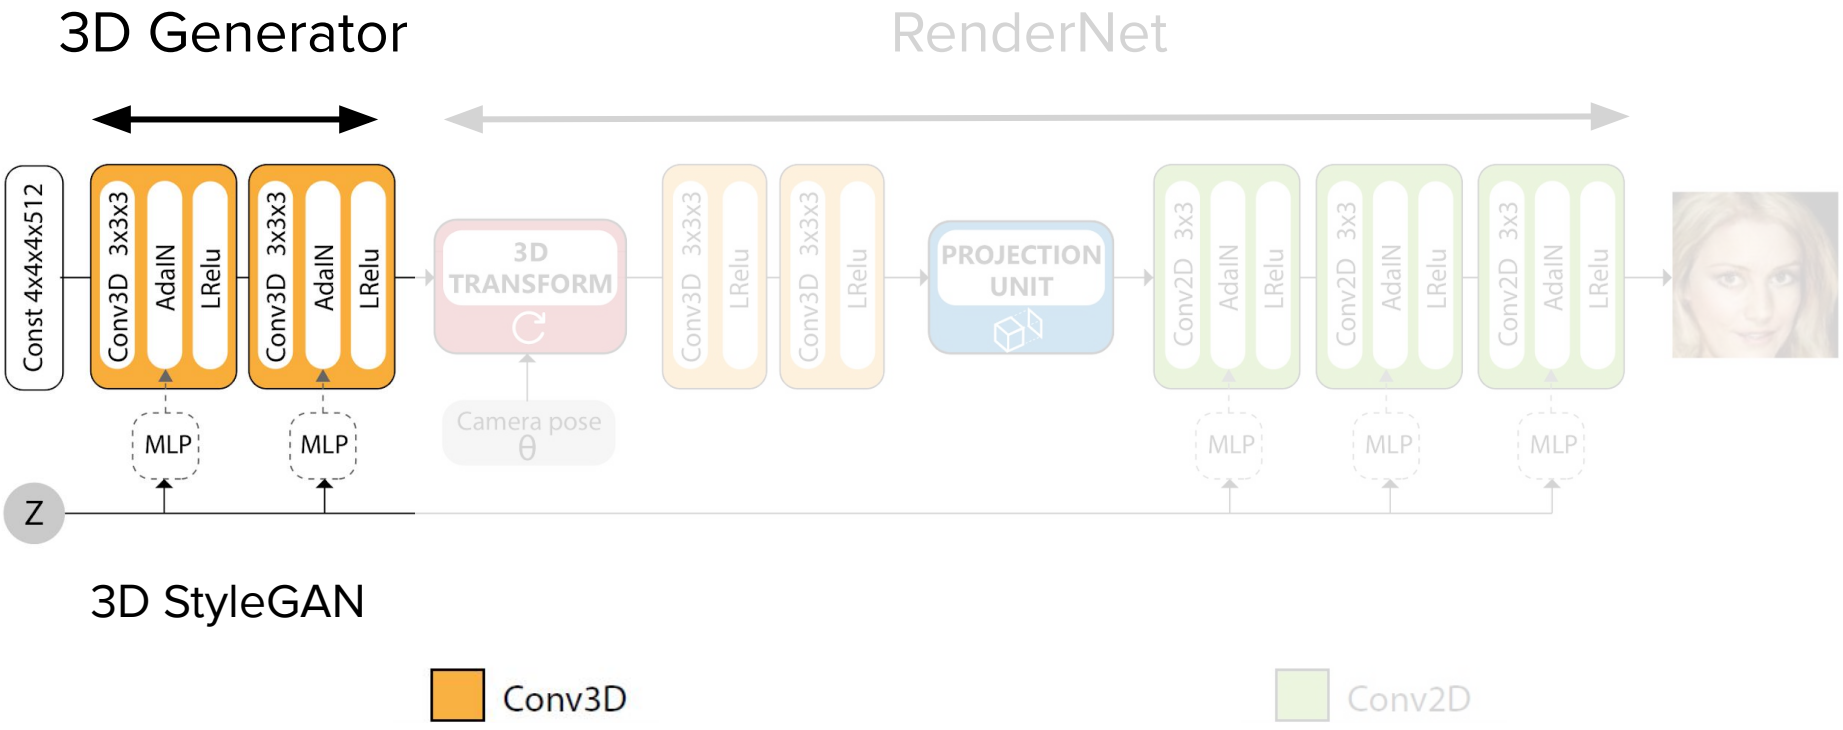

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

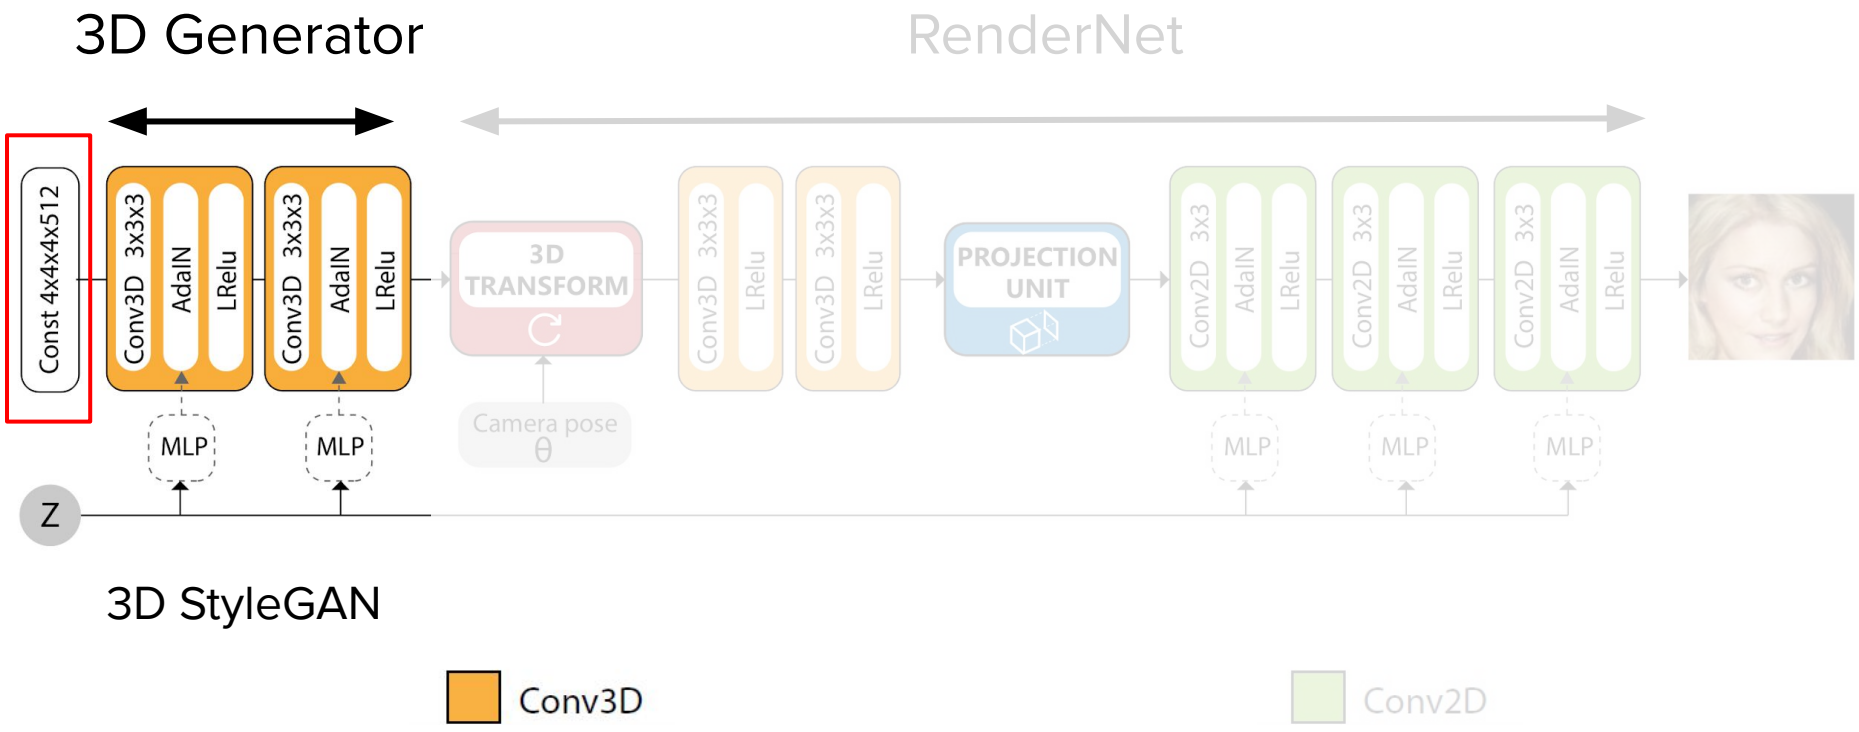

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

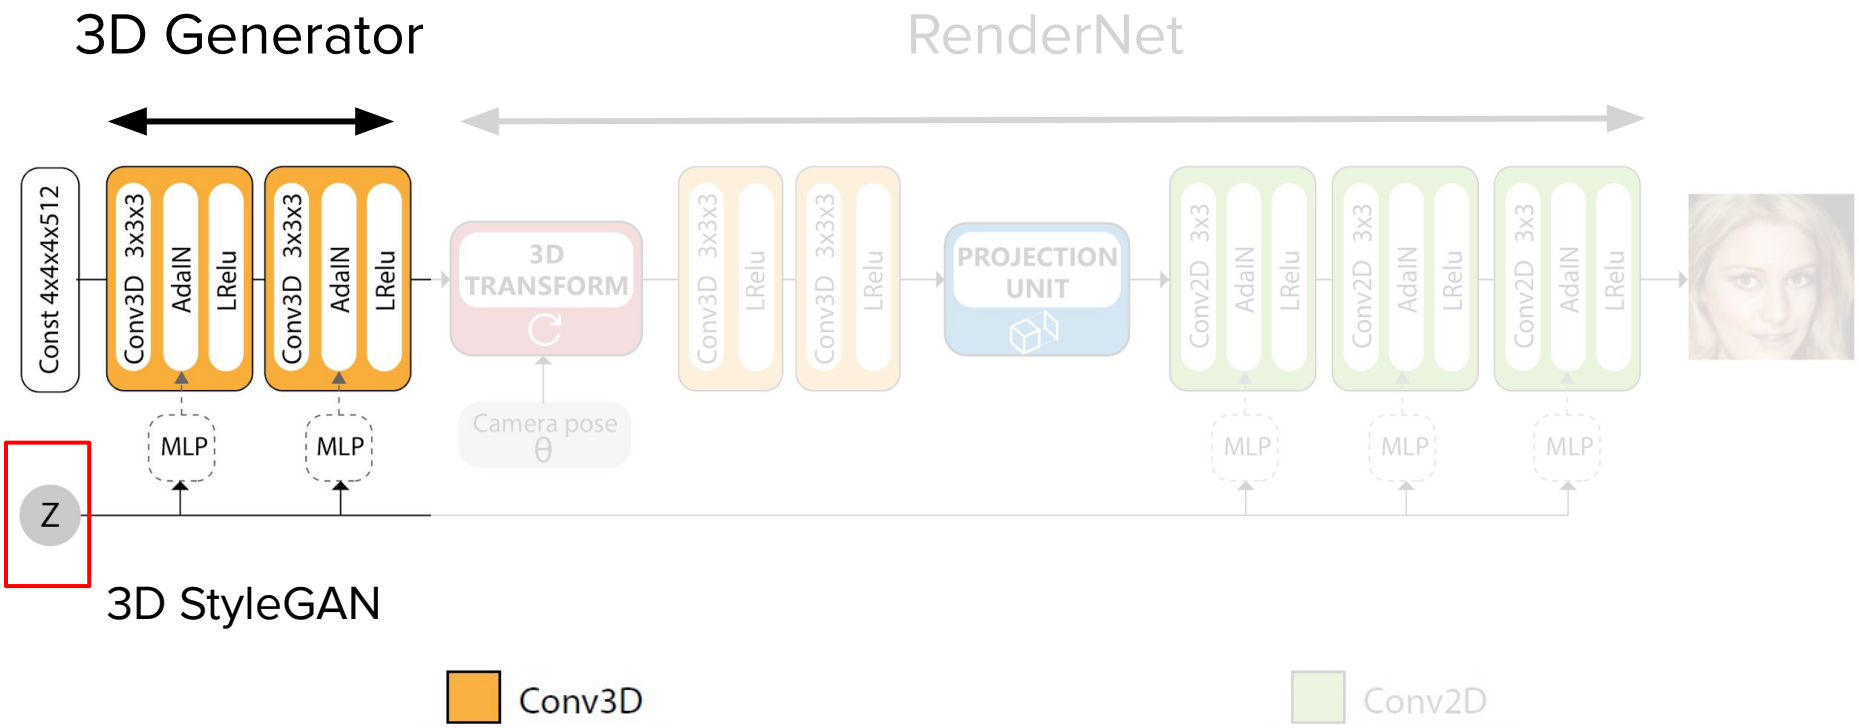

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

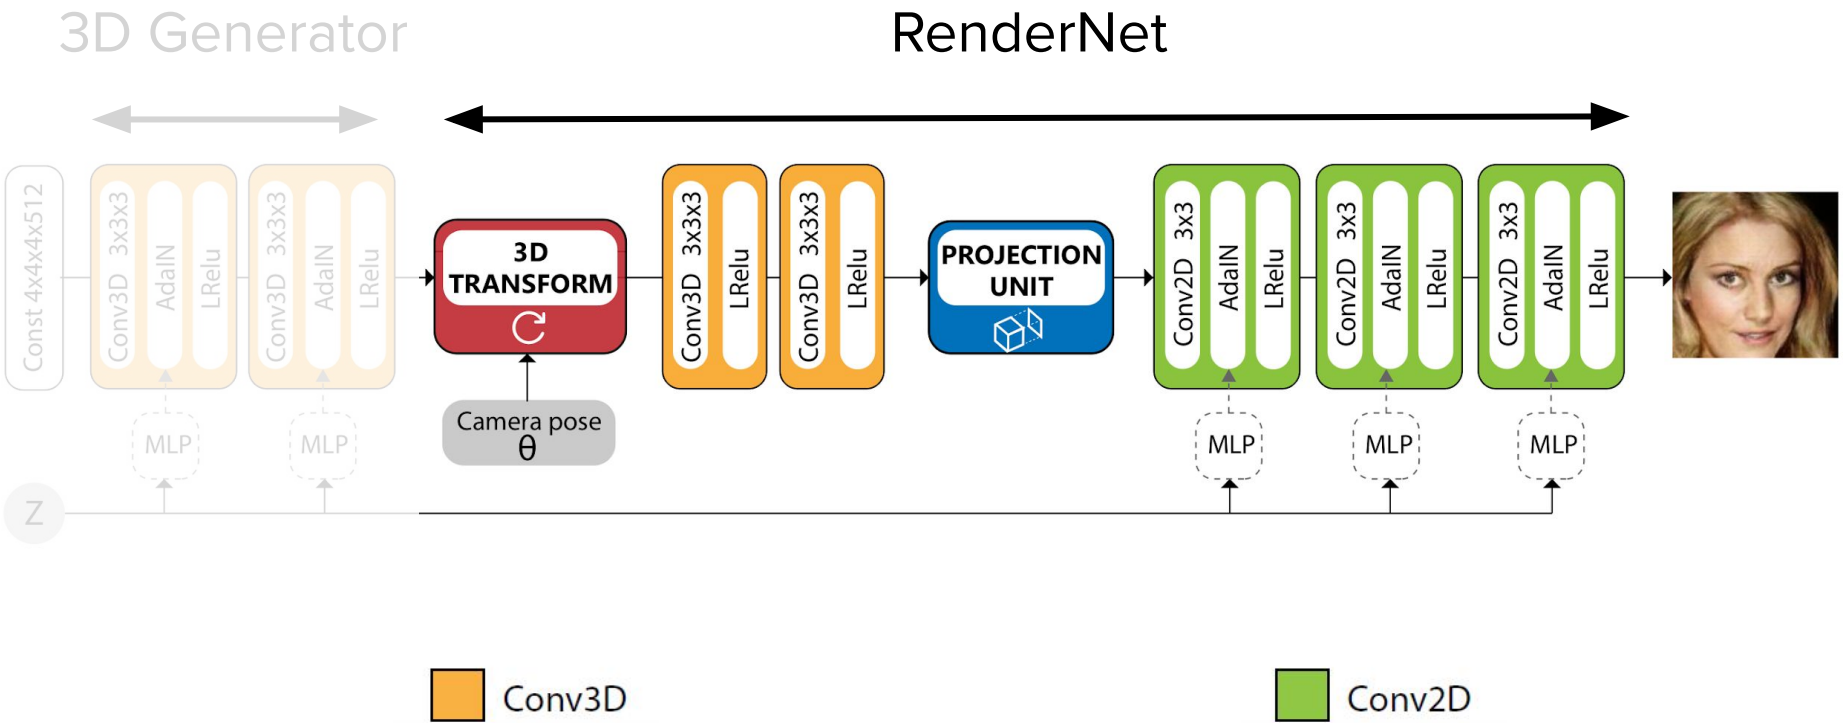

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

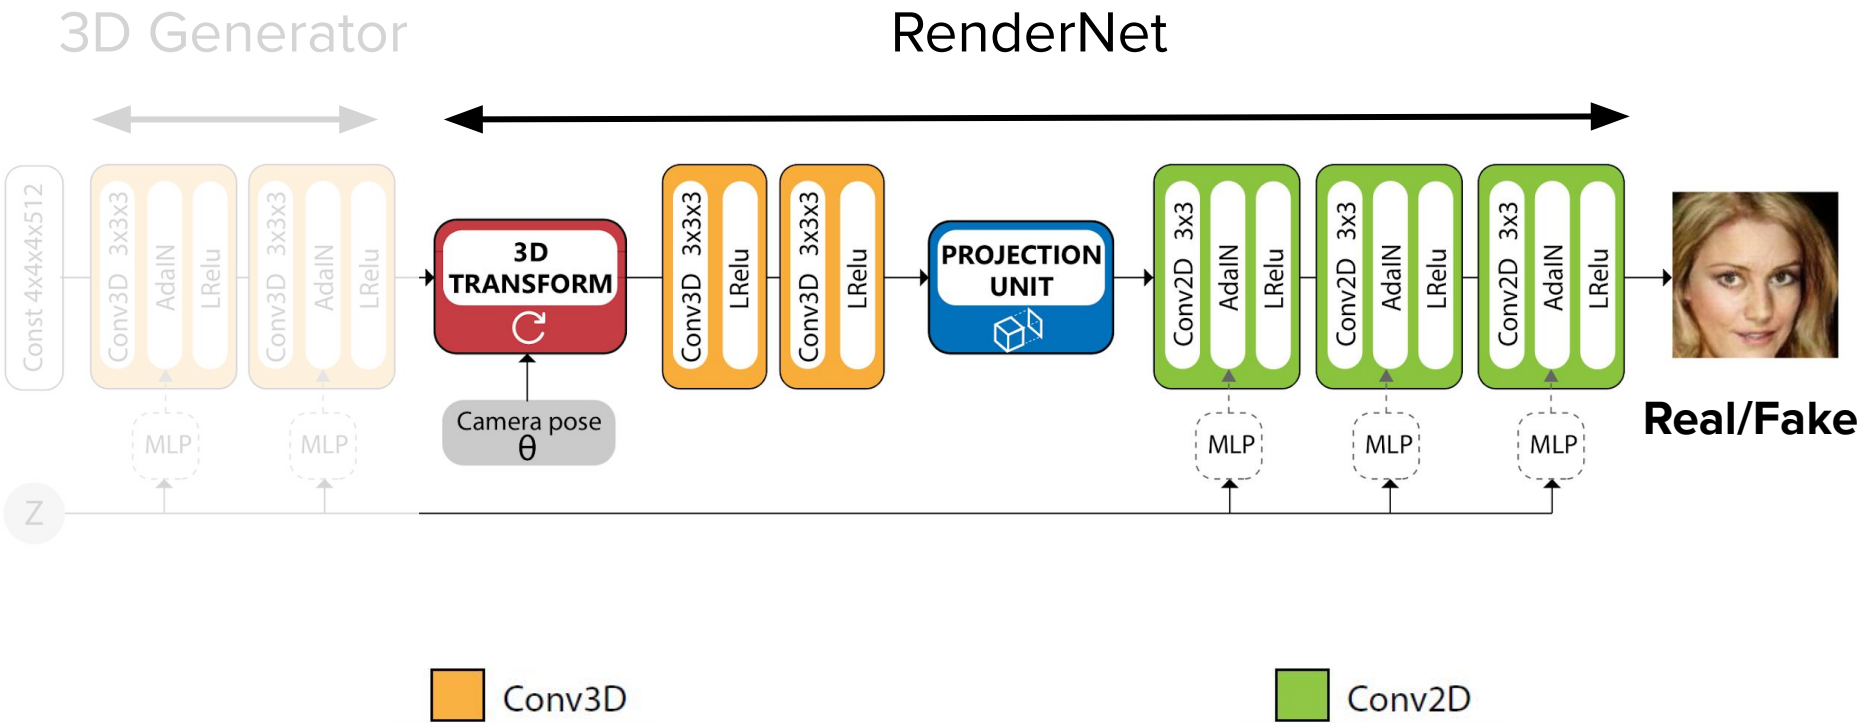

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

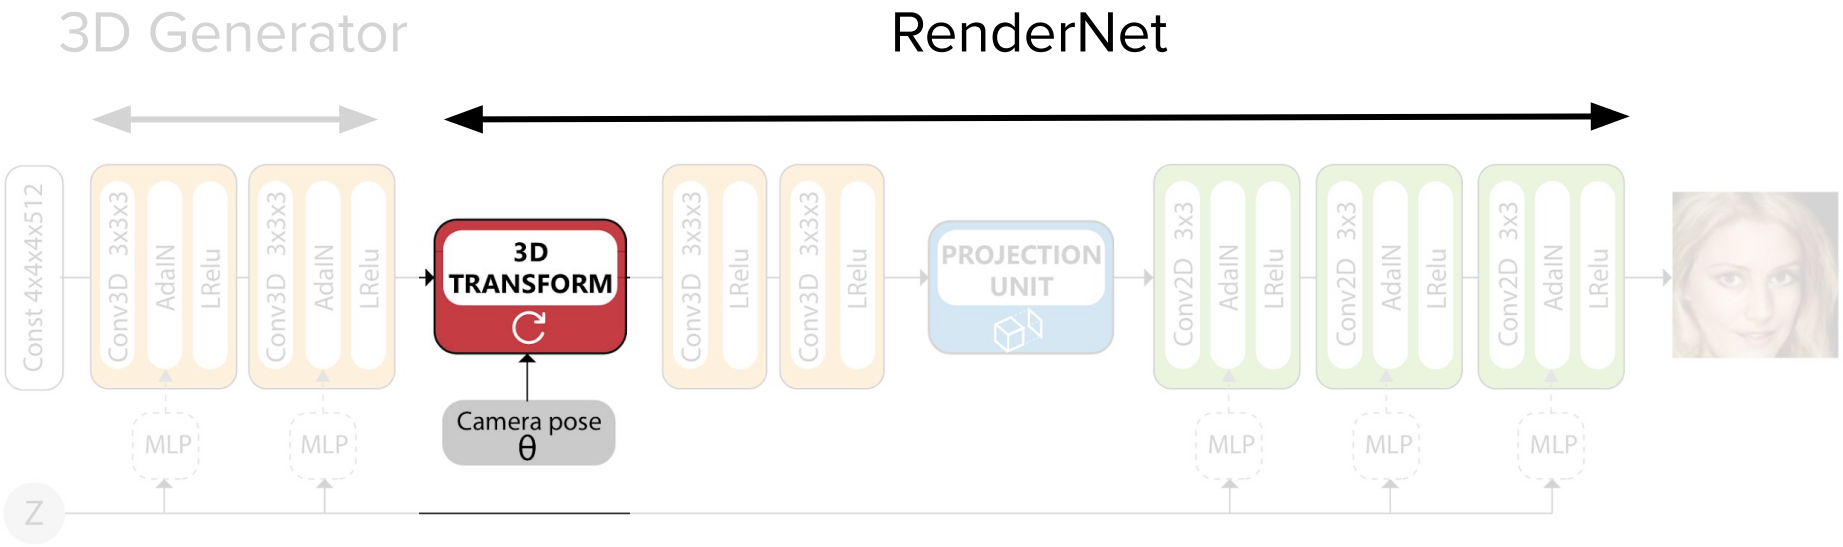

A representation that is unbreakable under 3D rigid-body transformations

Conv3D

Conv2D

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

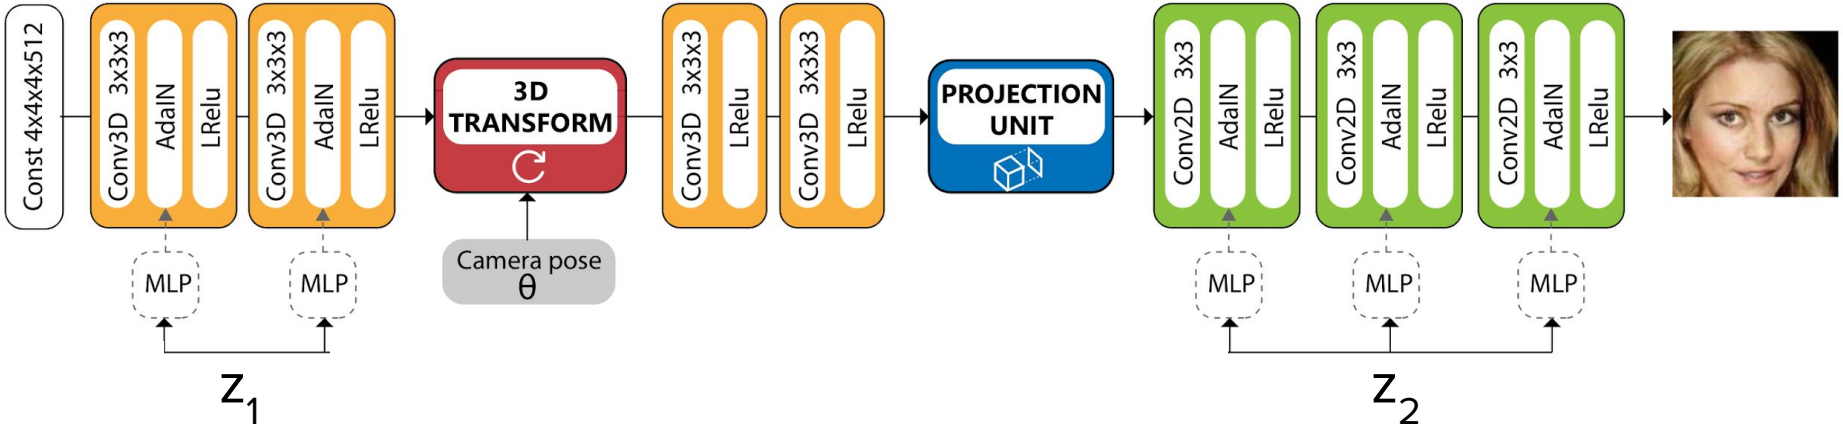

Shape Controller

Texture Controller

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

HoloGAN: Unsupervised learning of 3D representations from natural images

Thu Nguyen-Phuoc et al, ICCV 2019

Lambda

Forward (Computer Graphics)

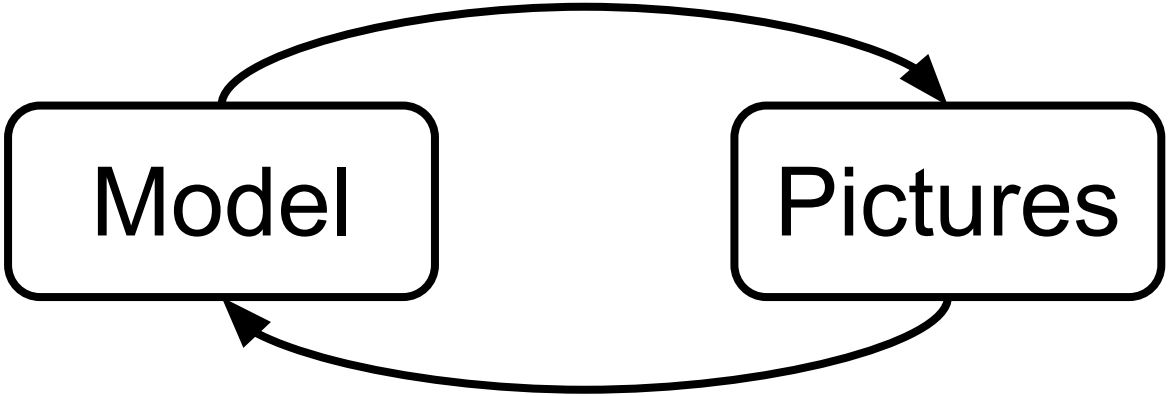

Inverse (Computer Vision)

Sub-module for Ray Tracing (Value / Policy Networks)

End-2-End Rasterization (Depthmap, Voxel, Point Cloud, Mesh)

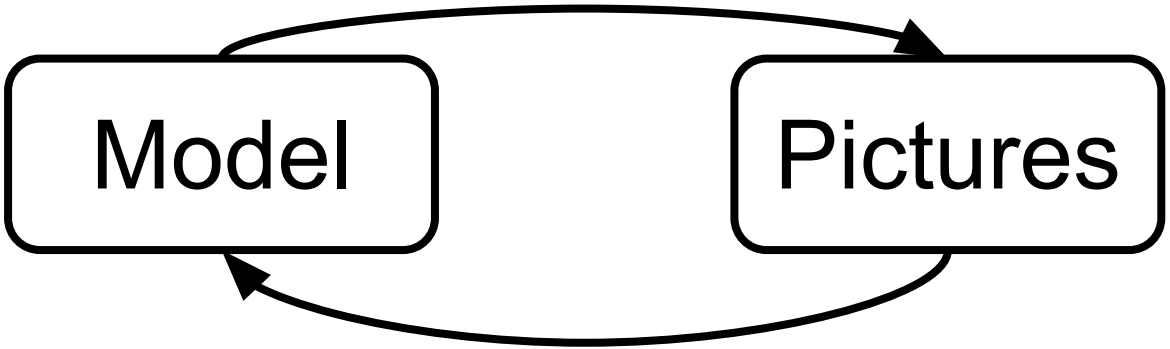

Differentiable Rendering (Representation Learning)

Thu Nguyen-Phuoc

Bing Xu

Yongliang Yang

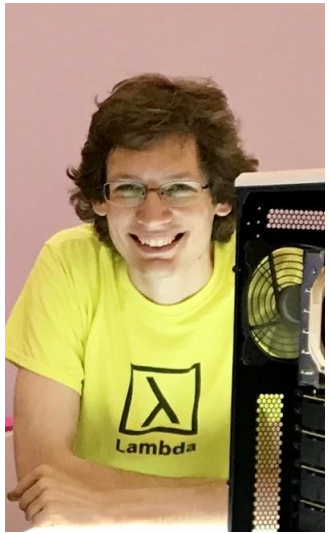

Stephen Balaban

Lucas Theis Christian Richardt Junfei Zhang Rui Wang Kun Xu Rui Tang

Lambda