Making a NeRF Relightable With Shadows

Xiuming Zhang 张修明

http://people.csail.mit.edu/xiuming/

GAMES 06/21/2021

Lighting is crucial to realism and aesthetics.

Caravaggio's "The Calling of Saint Matthew"

https://www.caravaggio-foundation.org/The-Calling-Of-Saint-Matthew.html

Paintings

Games

Lighting is crucial to realism and aesthetics.

Caravaggio's "The Calling of Saint Matthew"

https://www.caravaggio-foundation.org/The-Calling-Of-Saint-Matthew.html

Paintings

Games

"This Is the Way: How Innovative Technology Immersed Us in the World of the Mandalorian"

https://www.starwars.com/news/the-mandalorian-stagecraft-feature

Movies

Lighting is crucial to realism and aesthetics.

Caravaggio's "The Calling of Saint Matthew"

https://www.caravaggio-foundation.org/The-Calling-Of-Saint-Matthew.html

Paintings

Games

"This Is the Way: How Innovative Technology Immersed Us in the World of the Mandalorian"

https://www.starwars.com/news/the-mandalorian-stagecraft-feature

Movies

"Understanding Photography Lighting - The Essential Beginners Guide"

https://www.pixpa.com/blog/photography-lighting

Photography

Relighting has wide applications, …

Guo et al. [SIGGRAPH Asia 2019]

VR and AR

Relighting has wide applications, …

Guo et al. [SIGGRAPH Asia 2019]

VR and AR

Google's Portrait Light

Photo Enhancement

... but it is hard.

Local	Non-Local

... but it is hard.

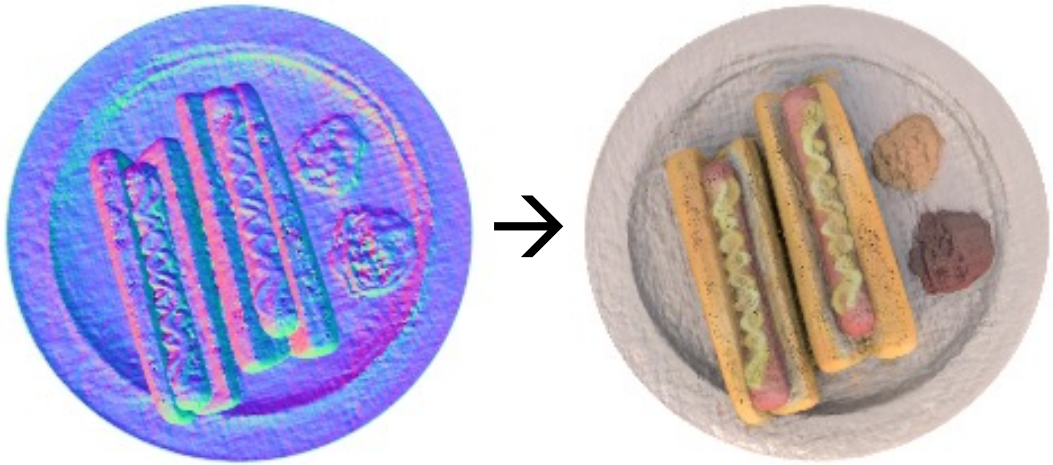

Geometry: bumpy normals → noisy render

Local | Non-Local

... but it is hard.

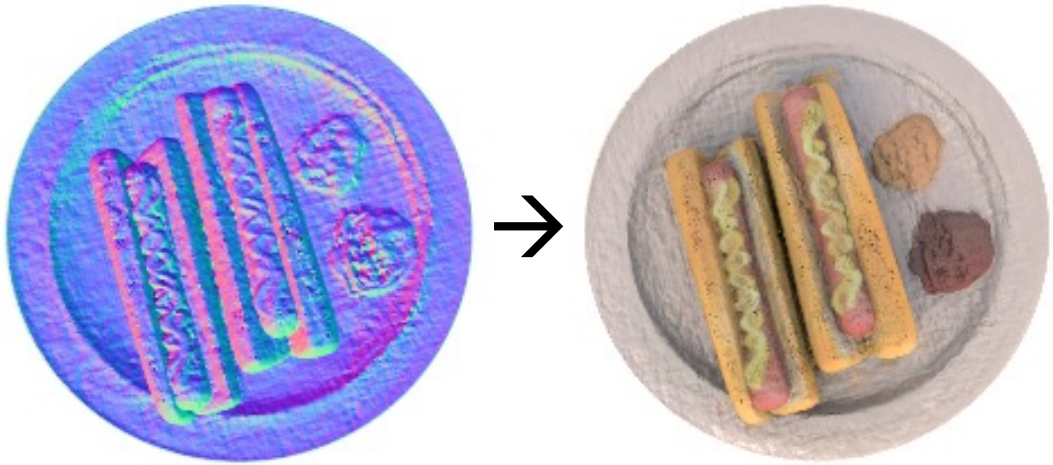

Geometry: bumpy normals → noisy render

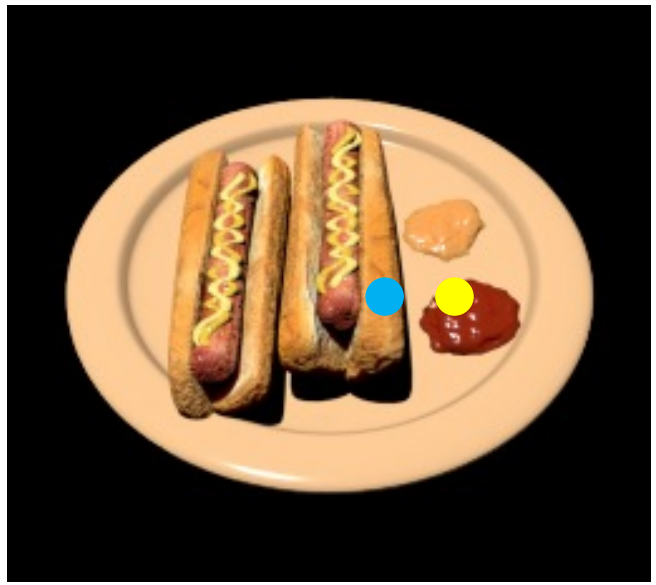

BRDFs: high-dimensional & spatially-varying

Local | Non-Local

... but it is hard.

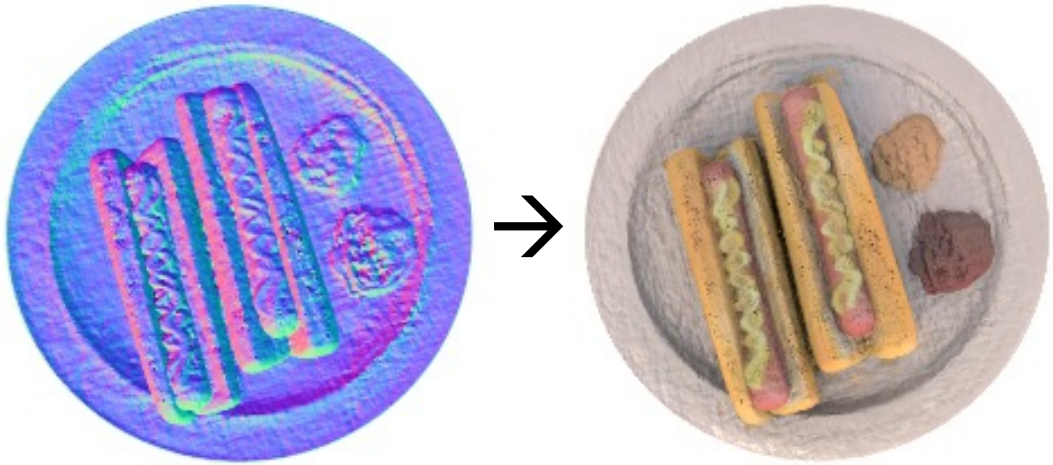

Geometry: bumpy normals → noisy render

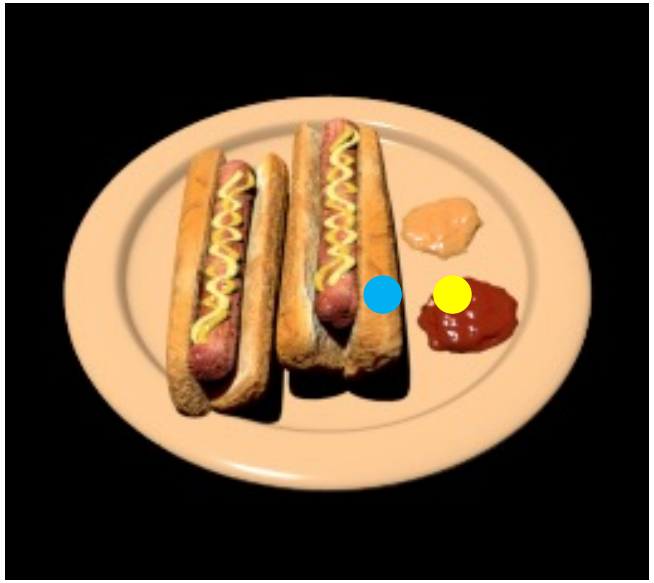

BRDFs: high-dimensional & spatially-varying

Shadowing: non-local geometry casting shadows

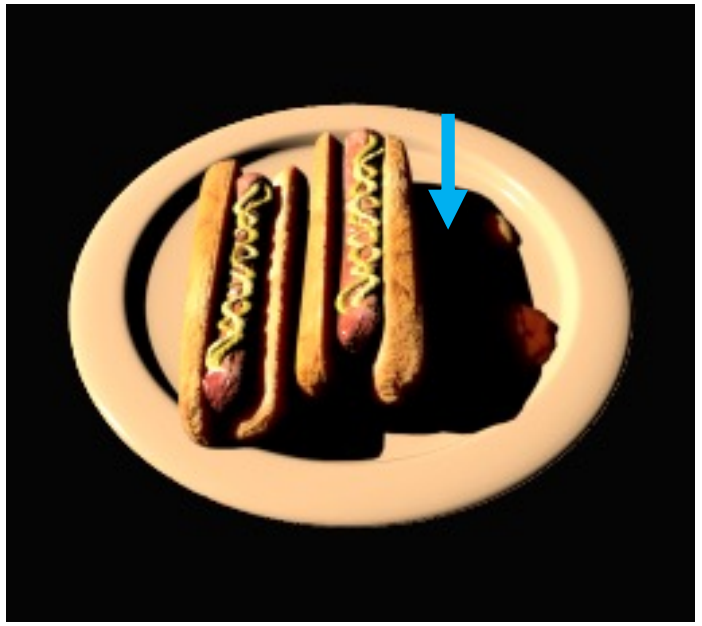

Local | Non-Local

... but it is hard.

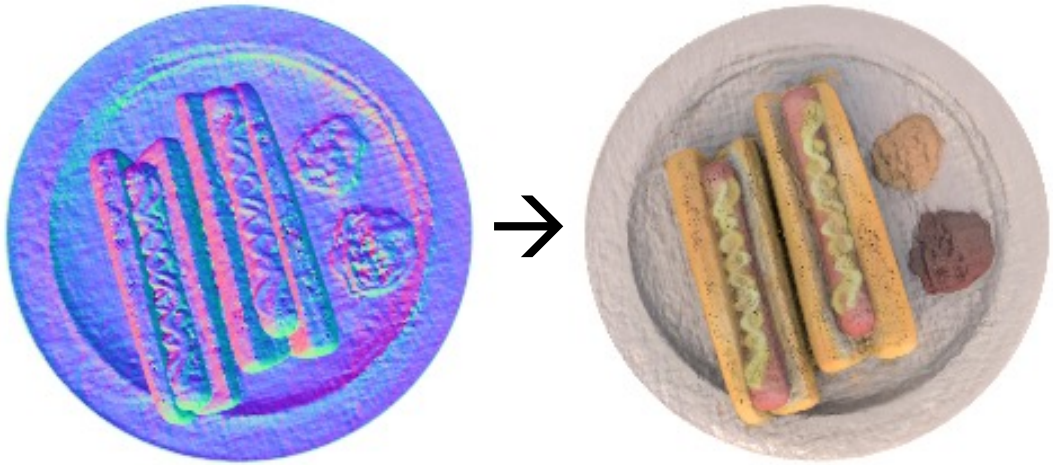

Geometry: bumpy normals → noisy render

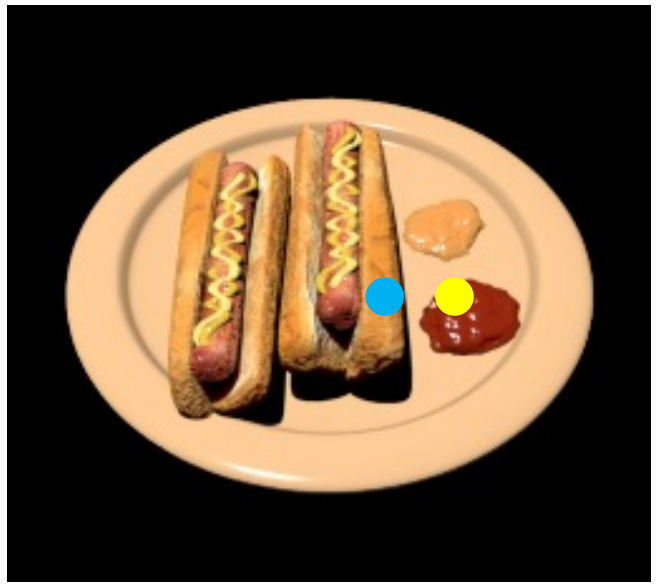

BRDFs: high-dimensional & spatially-varying

Shadowing: non-local geometry casting shadows

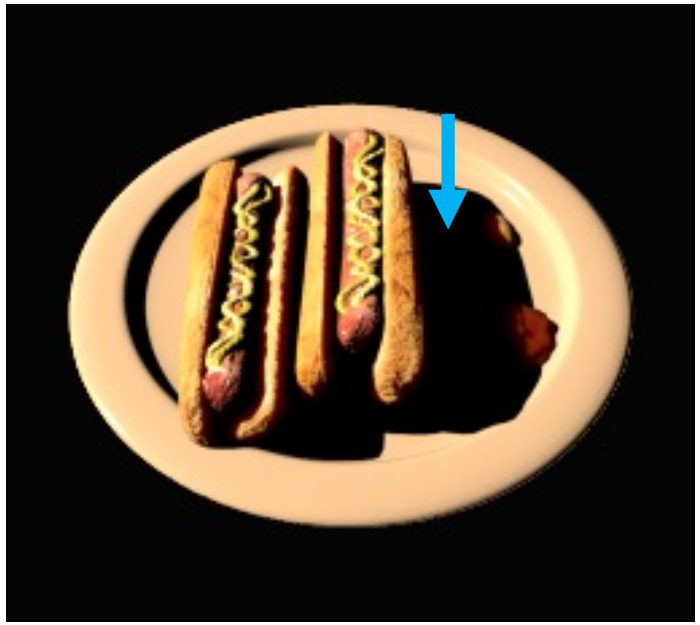

Global Illumination

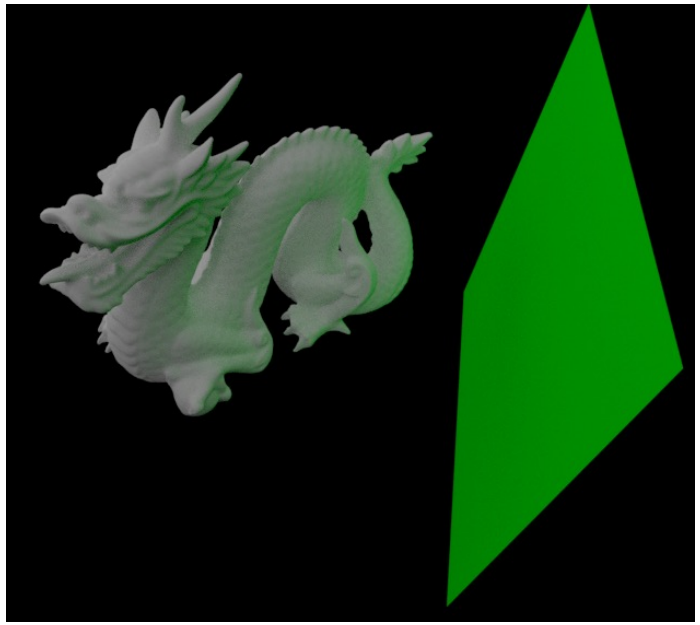

Local | Non-Local

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

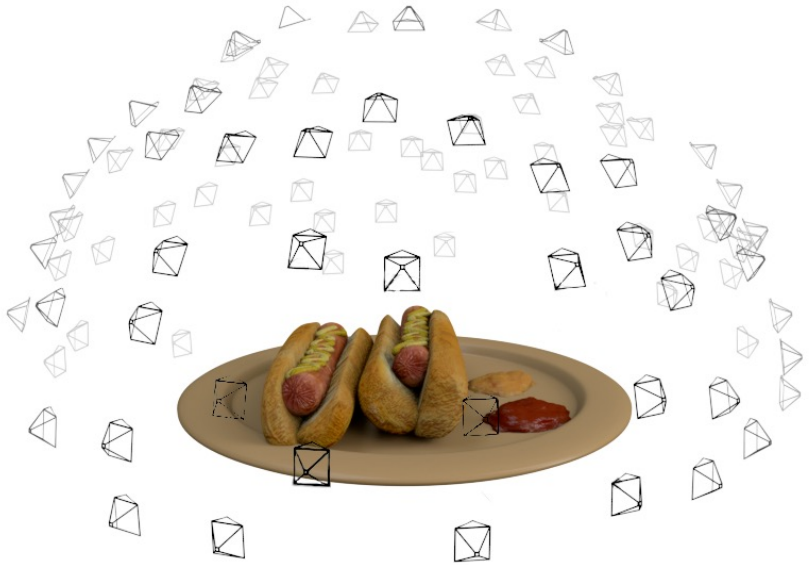

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

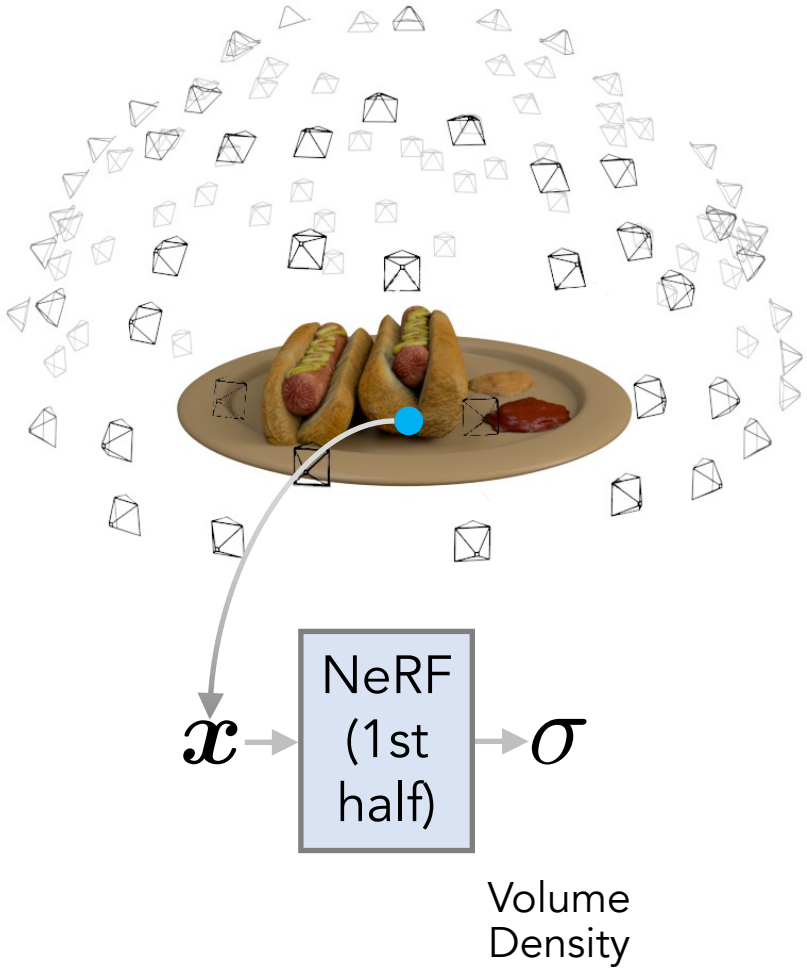

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

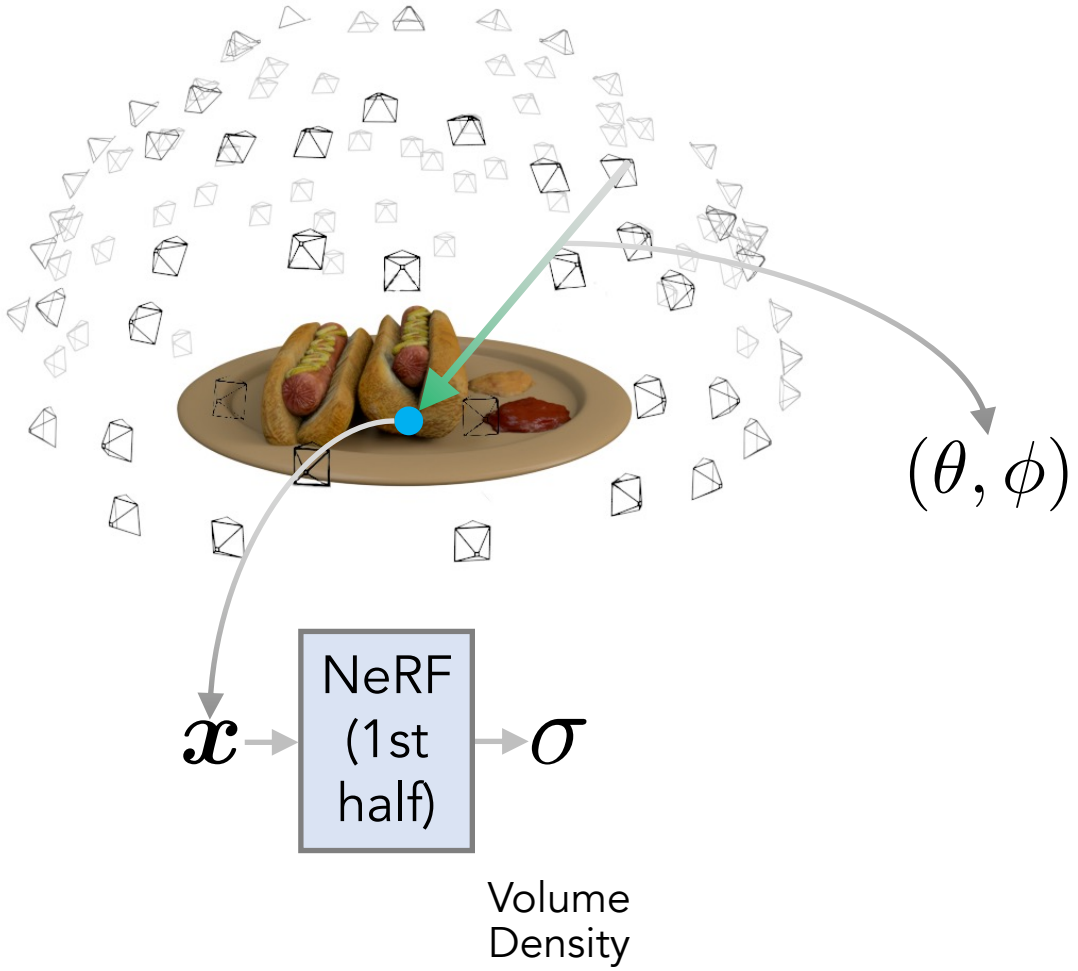

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

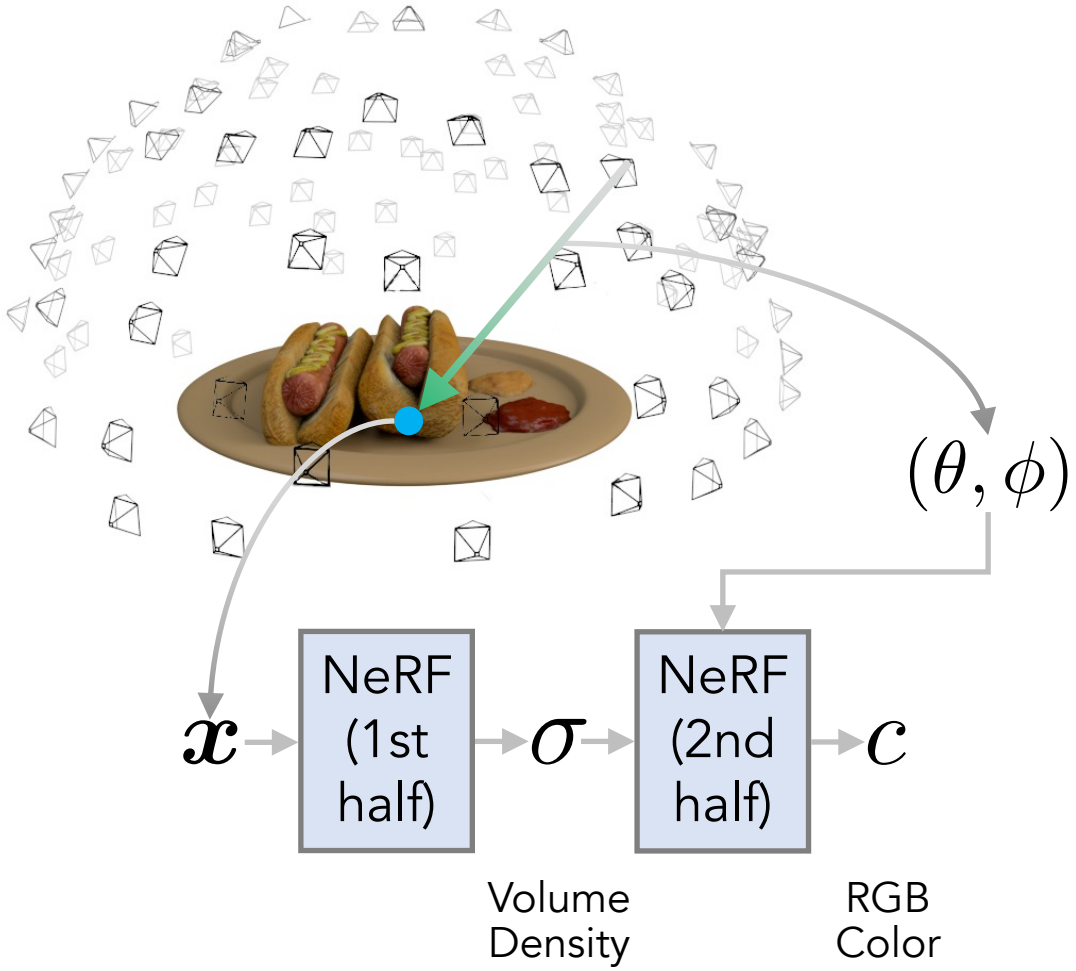

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

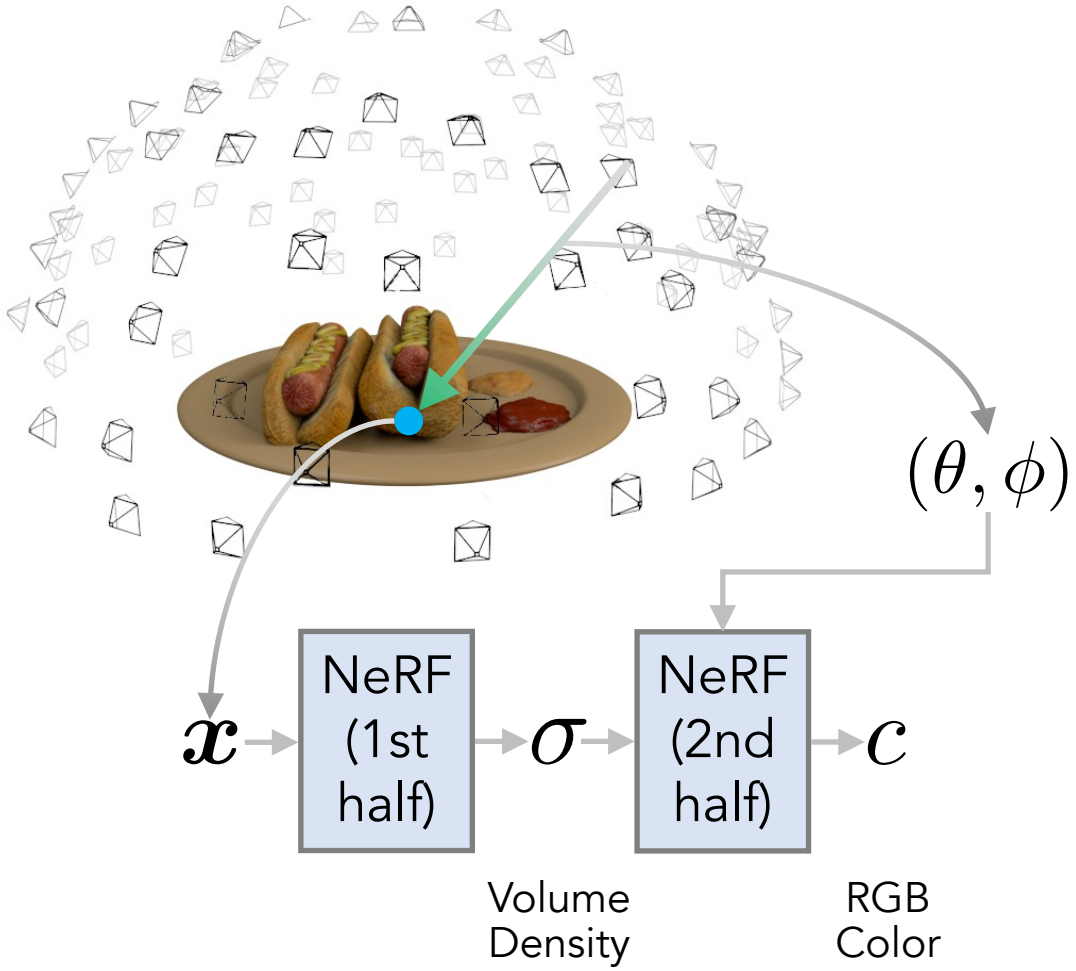

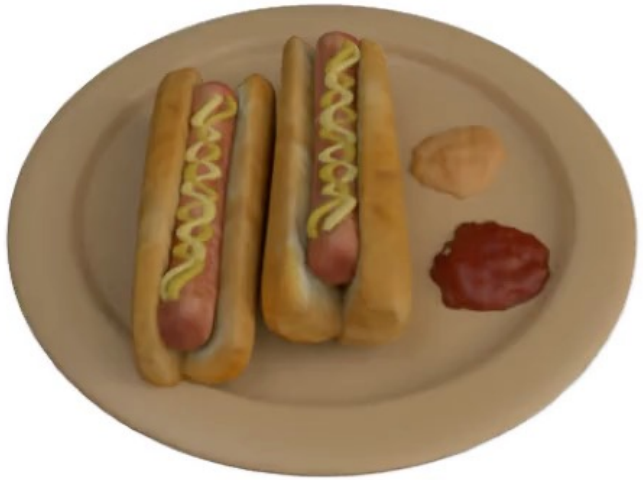

NeRF

✅ View Synthesis

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

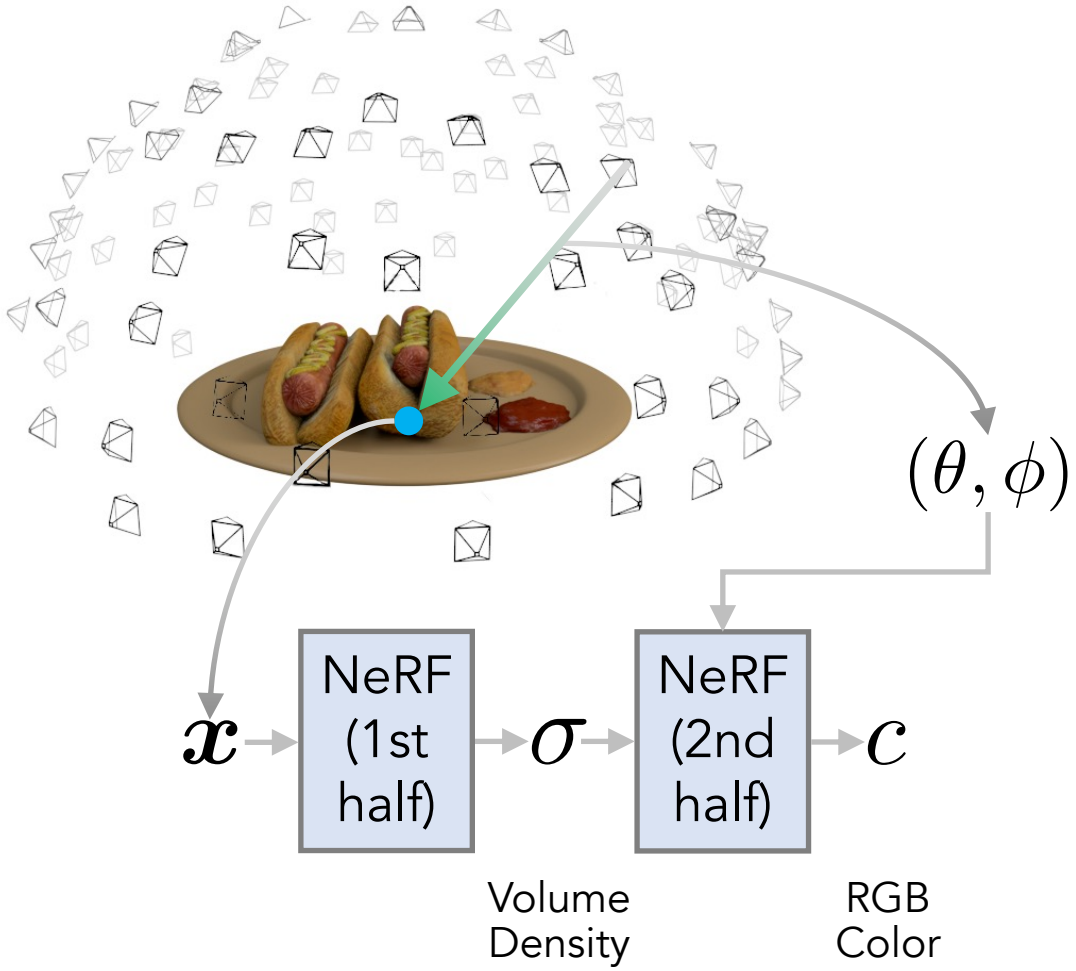

NeRF

✅ View Synthesis
❌ Relighting
❌ Material Editing

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

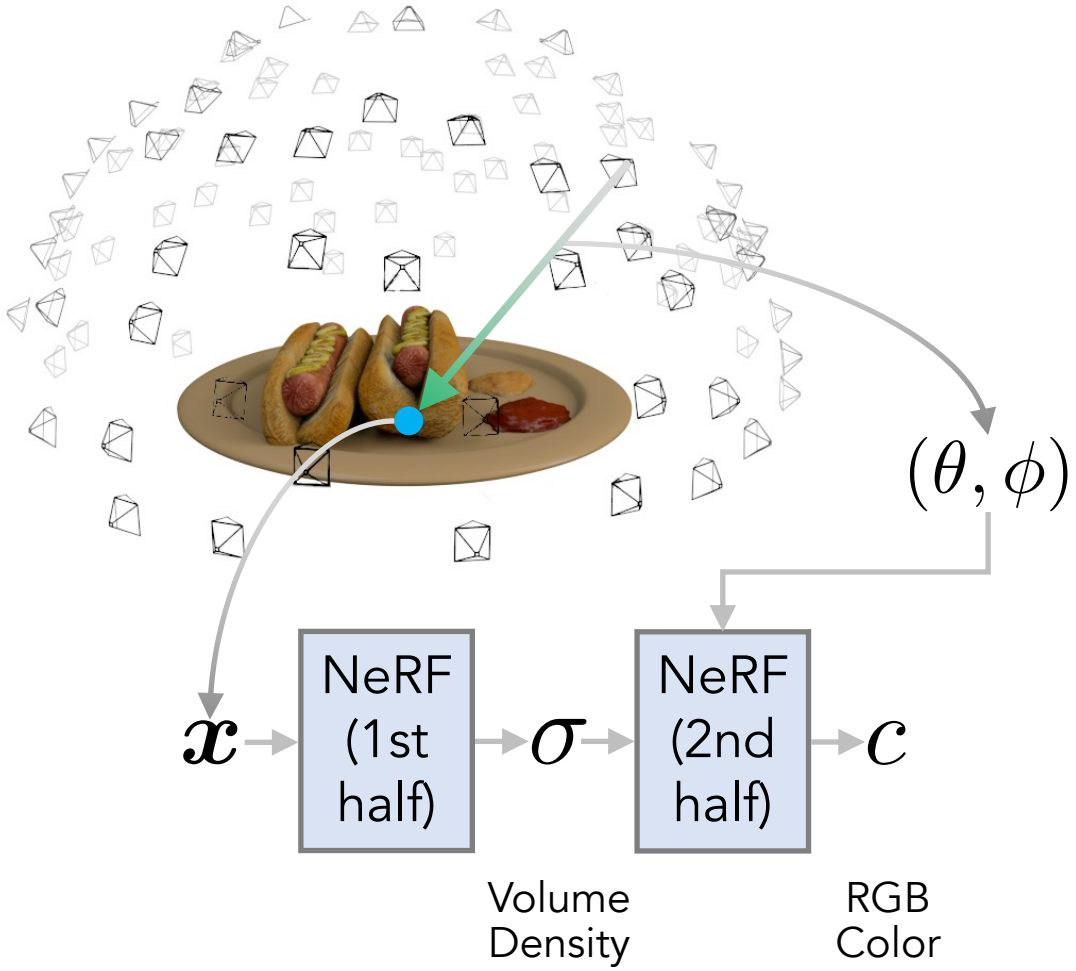

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Neural Radiance Fields (NeRF) [Mildenhall et al. 2020]

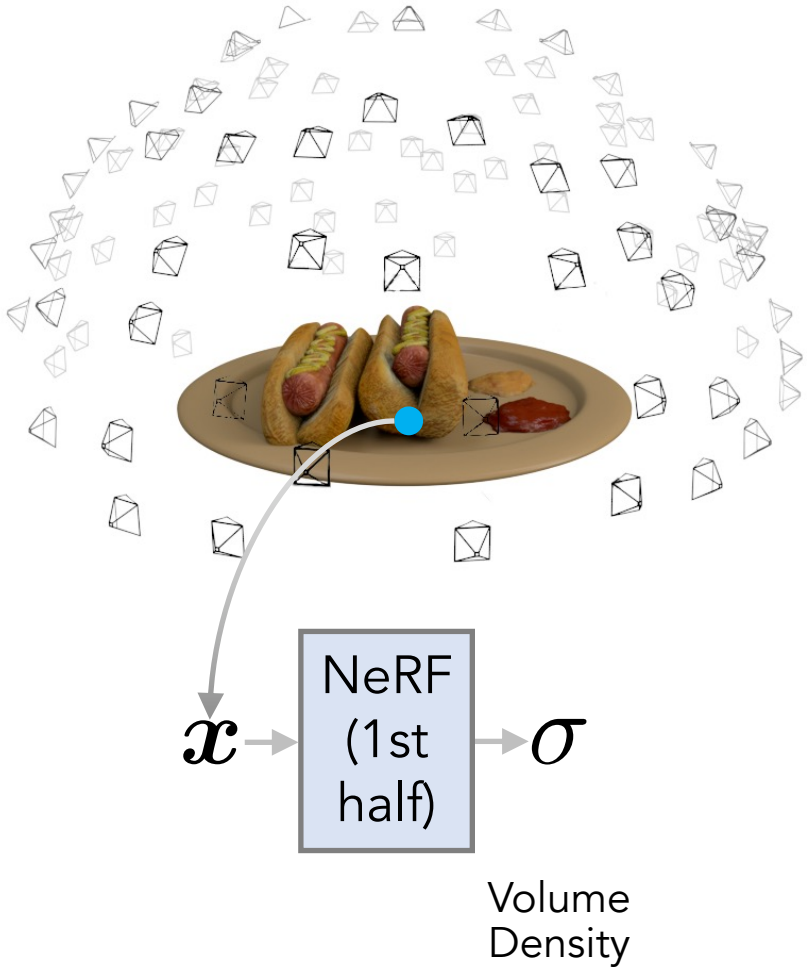

Ben Mildenhall*, Pratul P. Srinivasan*, Matthew Tancik*, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng
NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
ECCV 2020

Normals

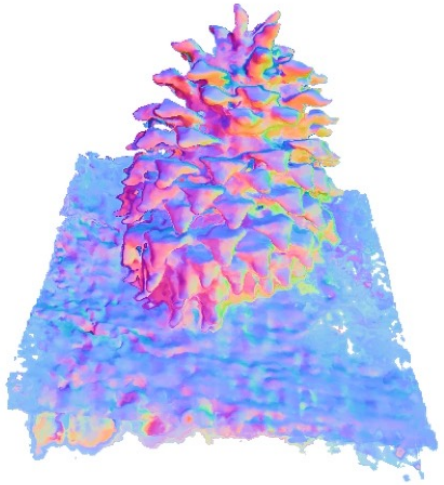

Visibility

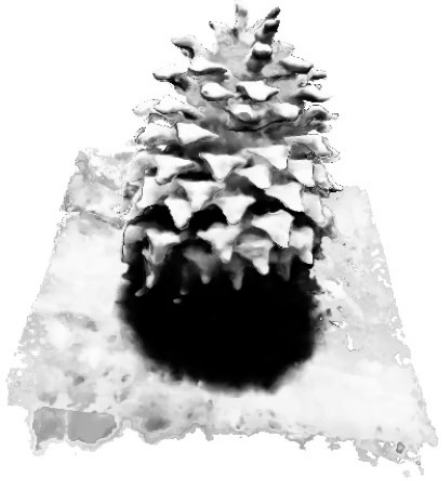

BRDF

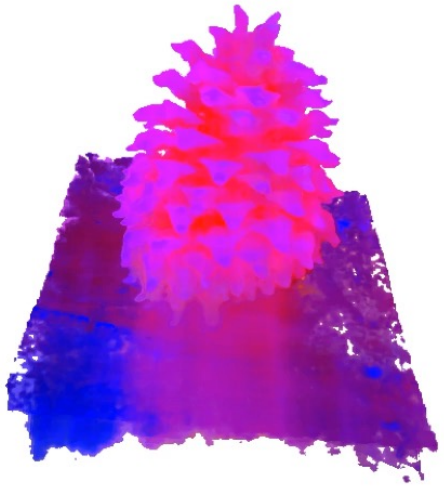

Albedo

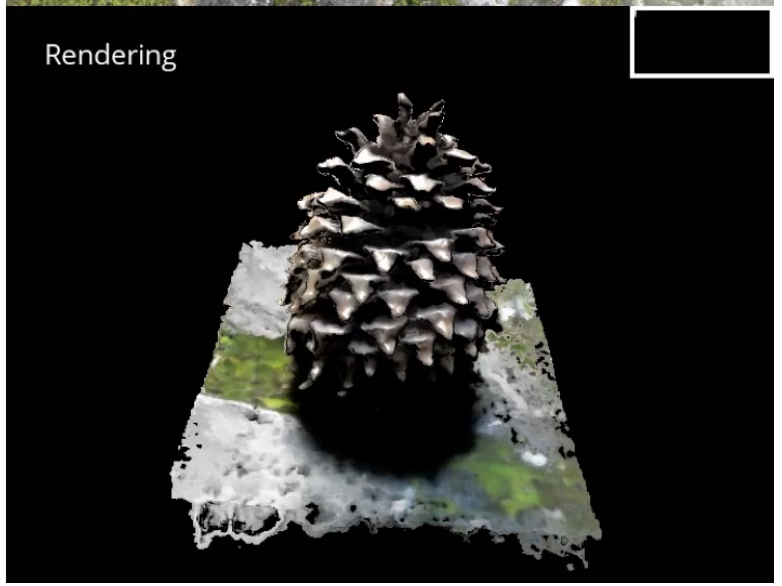

NeRFactor

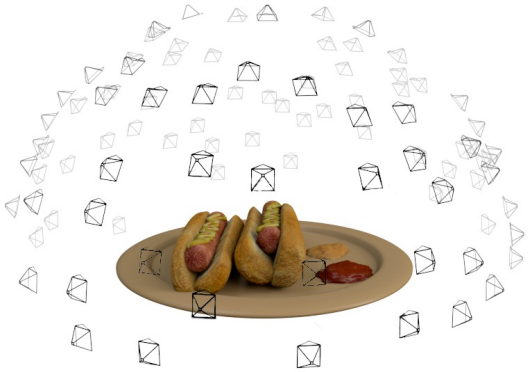

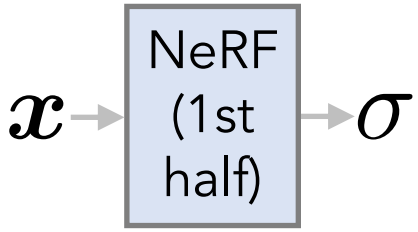

Volume Density

Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, Jonathan T. Barron
NeRFactor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination
arXiv

NeRFactor

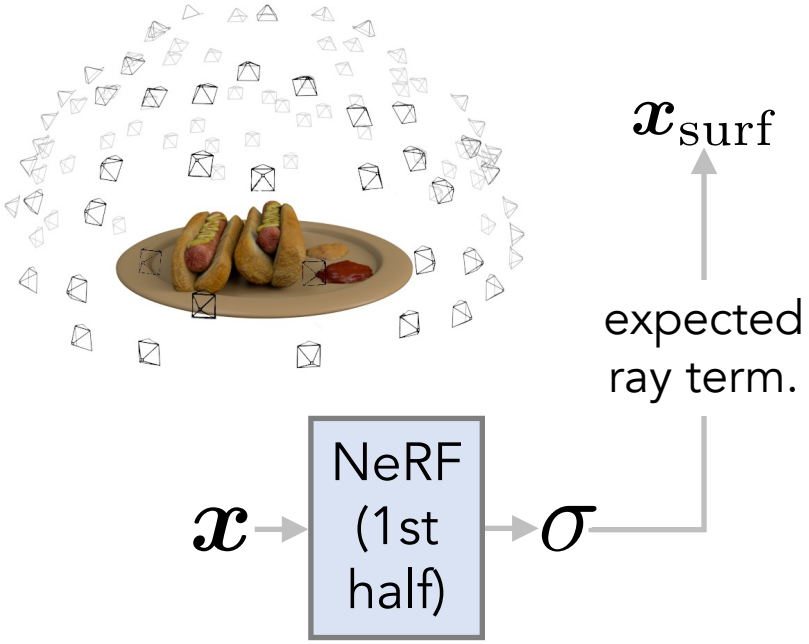
$\boldsymbol{x}_{\text{surf}}$
expected
ray term.
$\boldsymbol{x}$
NeRF
(1st
half)
σ

NeRFactor

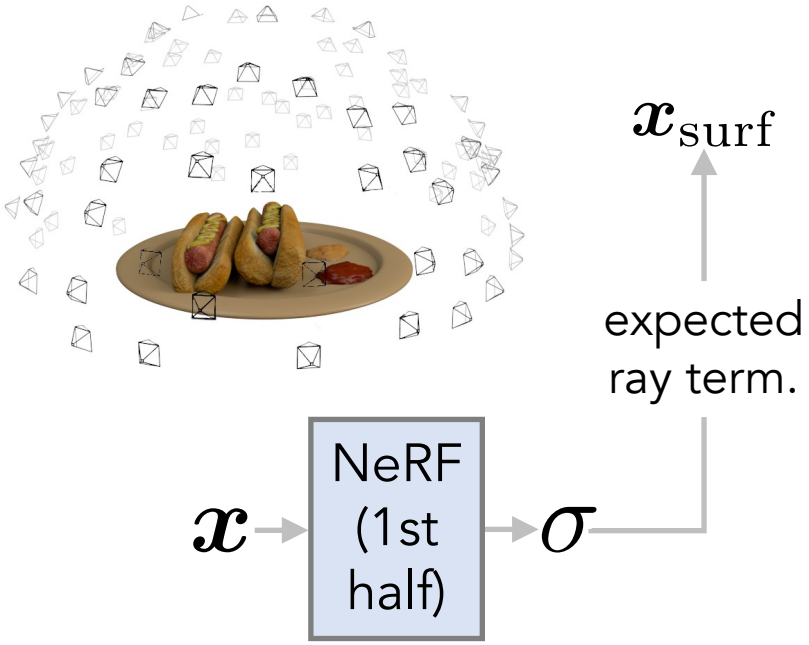

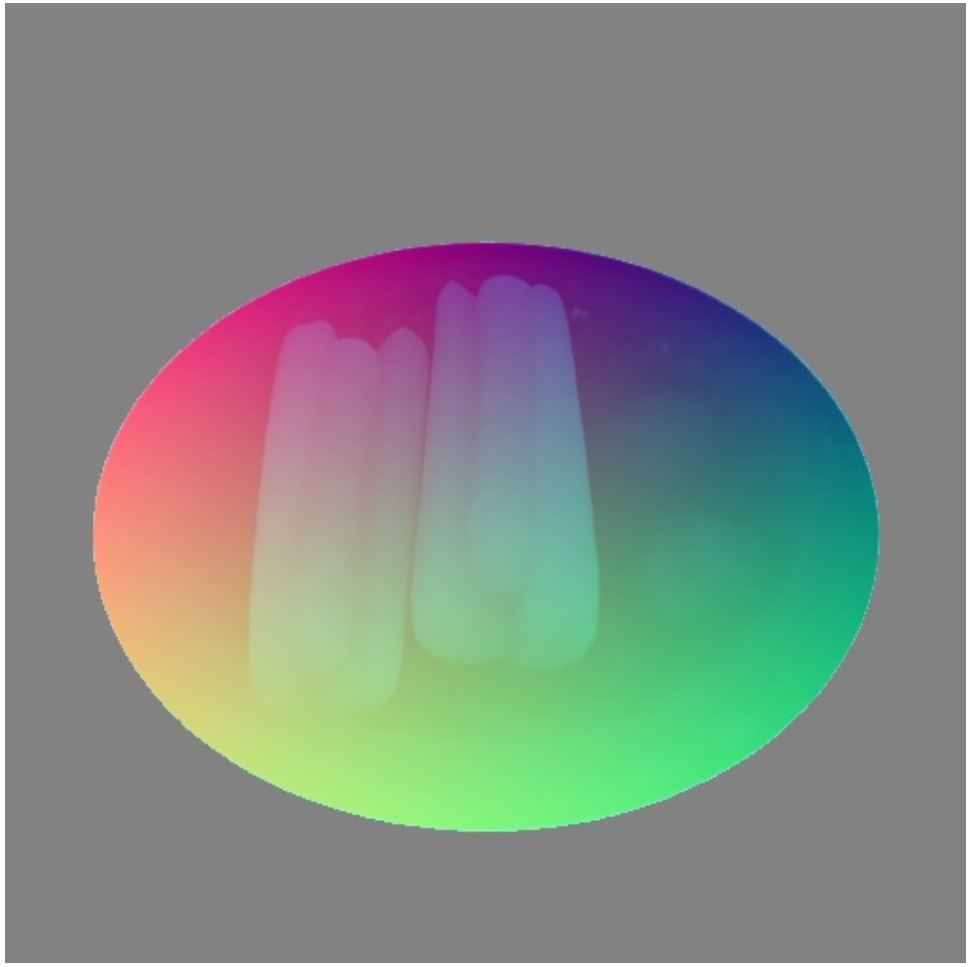

NeRF's Surface XYZs

NeRFactor

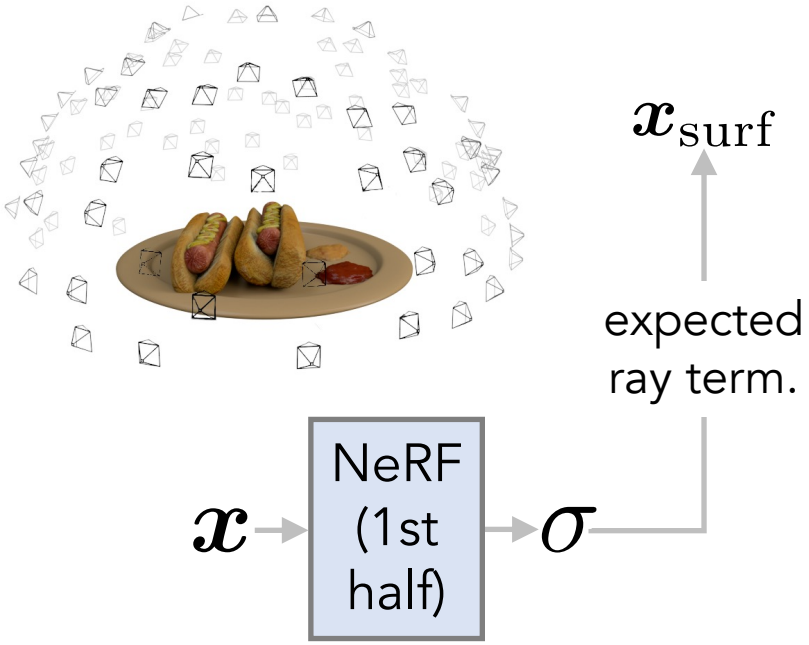
$\boldsymbol{x}_{\text{surf}}$
expected
ray term.
$\boldsymbol{x}$
NeRF
(1st
half)
σ

NeRFactor

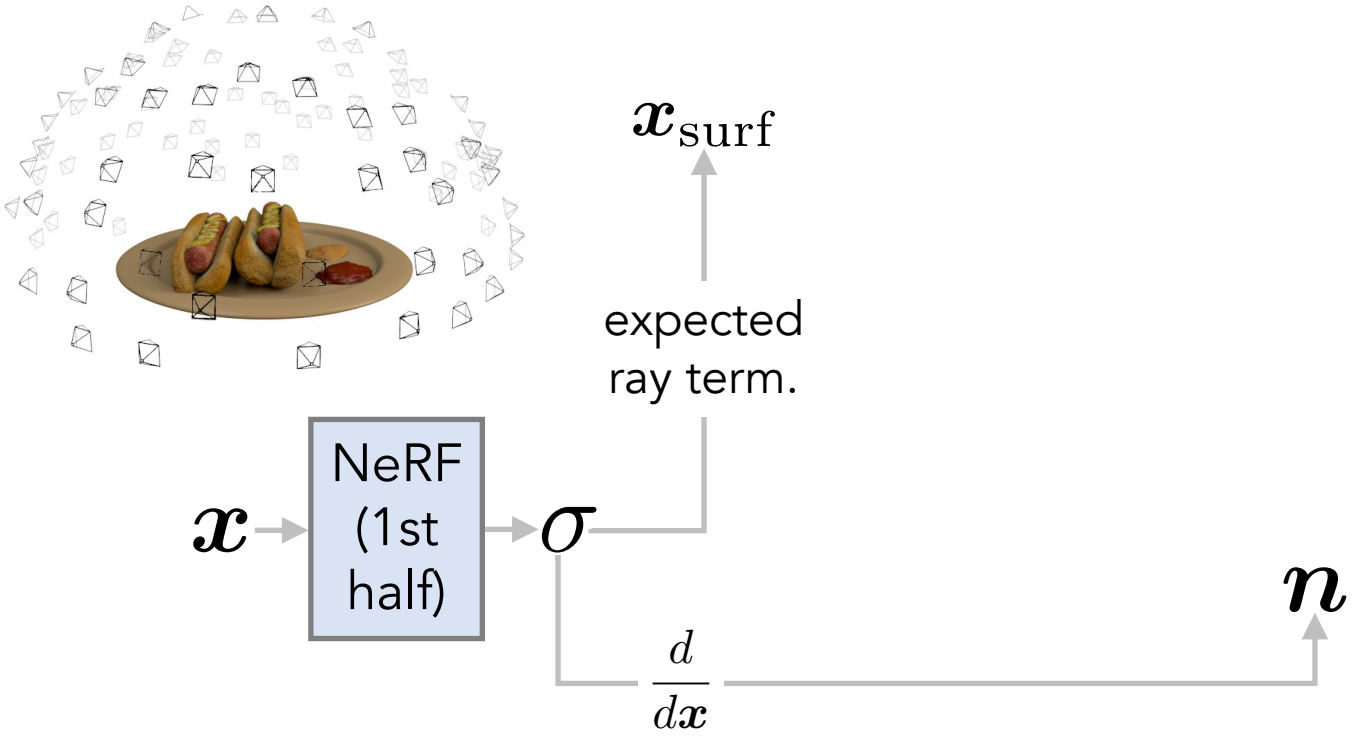
$\boldsymbol{x}_{\text{surf}}$
expected
ray term.
$\boldsymbol{x}$
NeRF
(1st
half)
σ
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$

NeRFactor

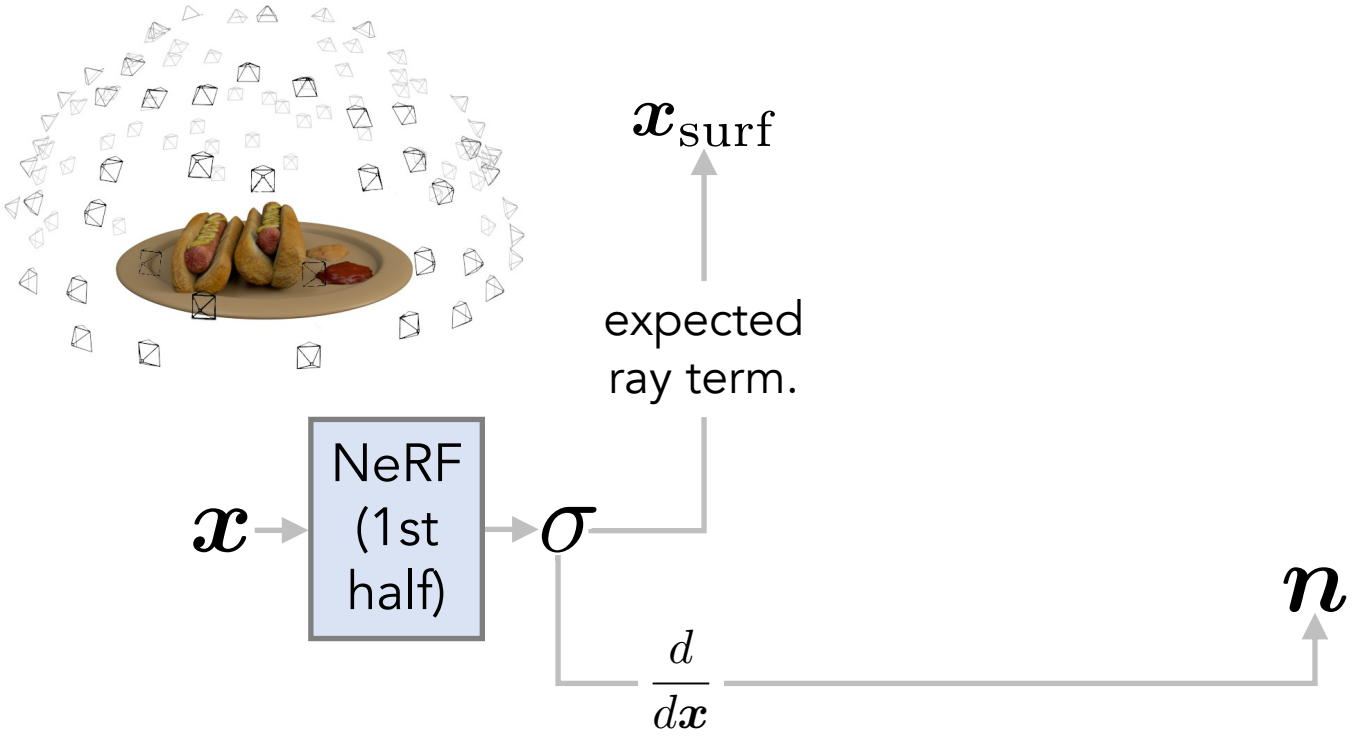

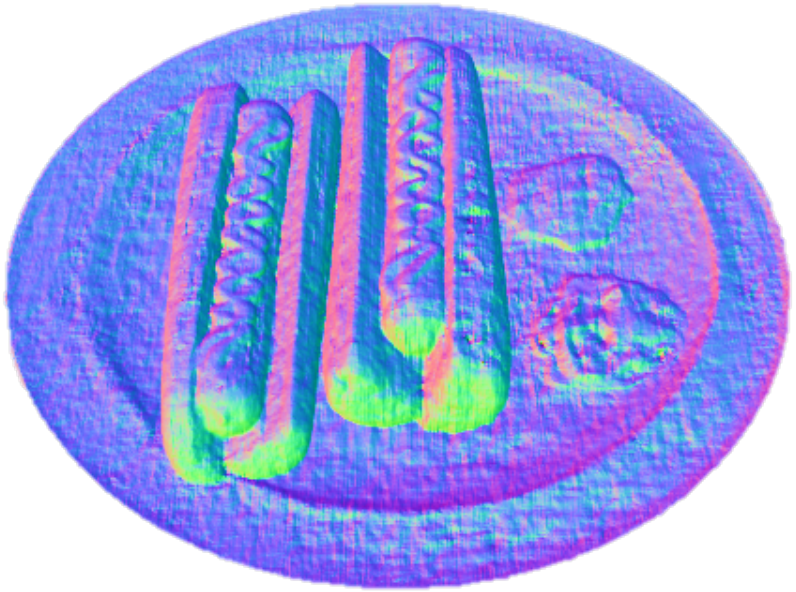

NeRF's Surface Normals

NeRFactor

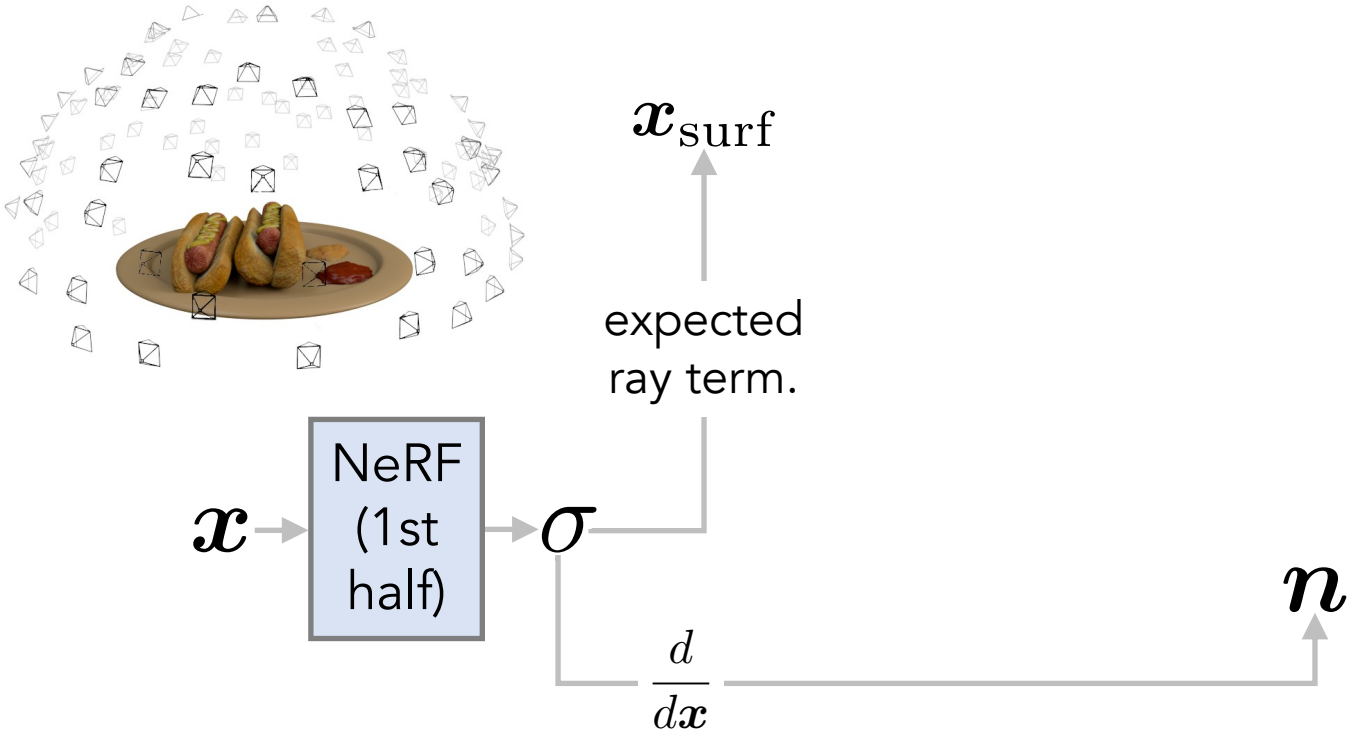
$\boldsymbol{x}_{\text{surf}}$
expected
ray term.
$\boldsymbol{x}$
NeRF
(1st
half)
σ
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$

NeRFactor

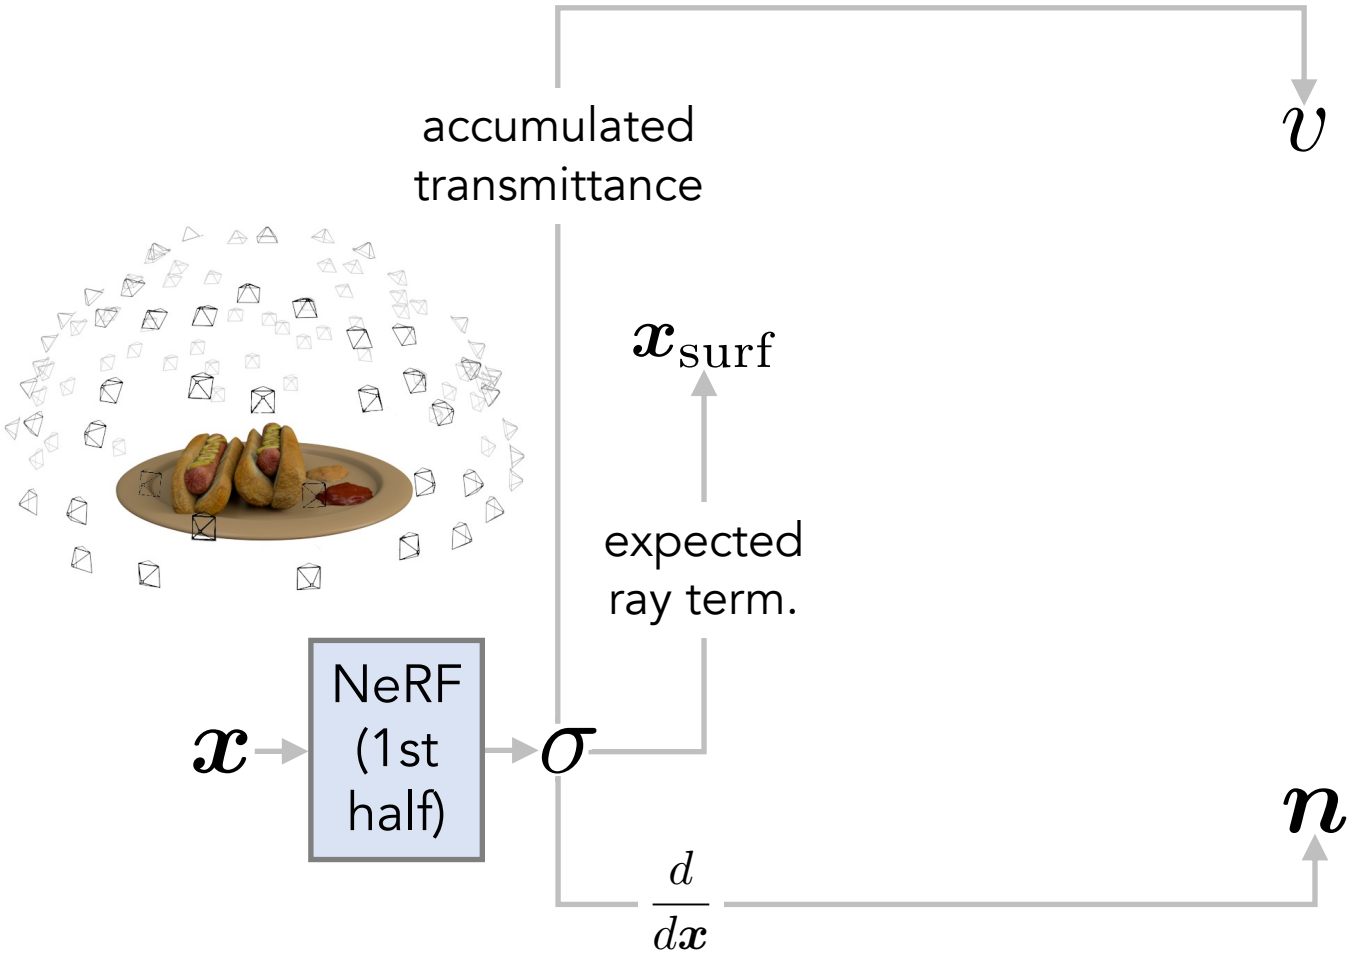
accumulated
transmittance
v
$\boldsymbol{x}_{\text{surf}}$
expected
ray term.
$\boldsymbol{x}$
NeRF
(1st
half)
σ
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$

NeRFactor

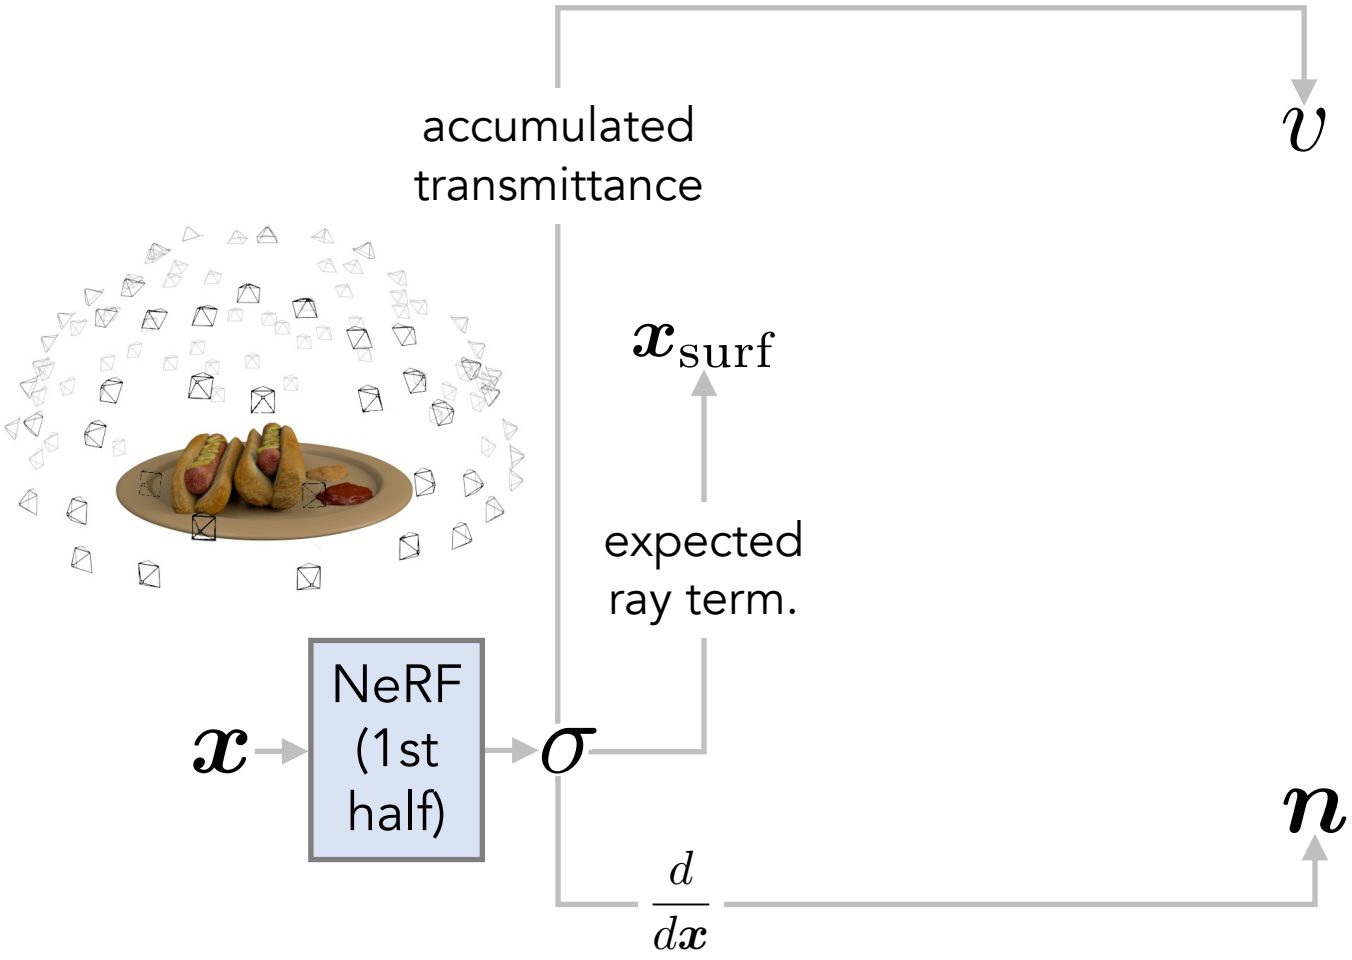

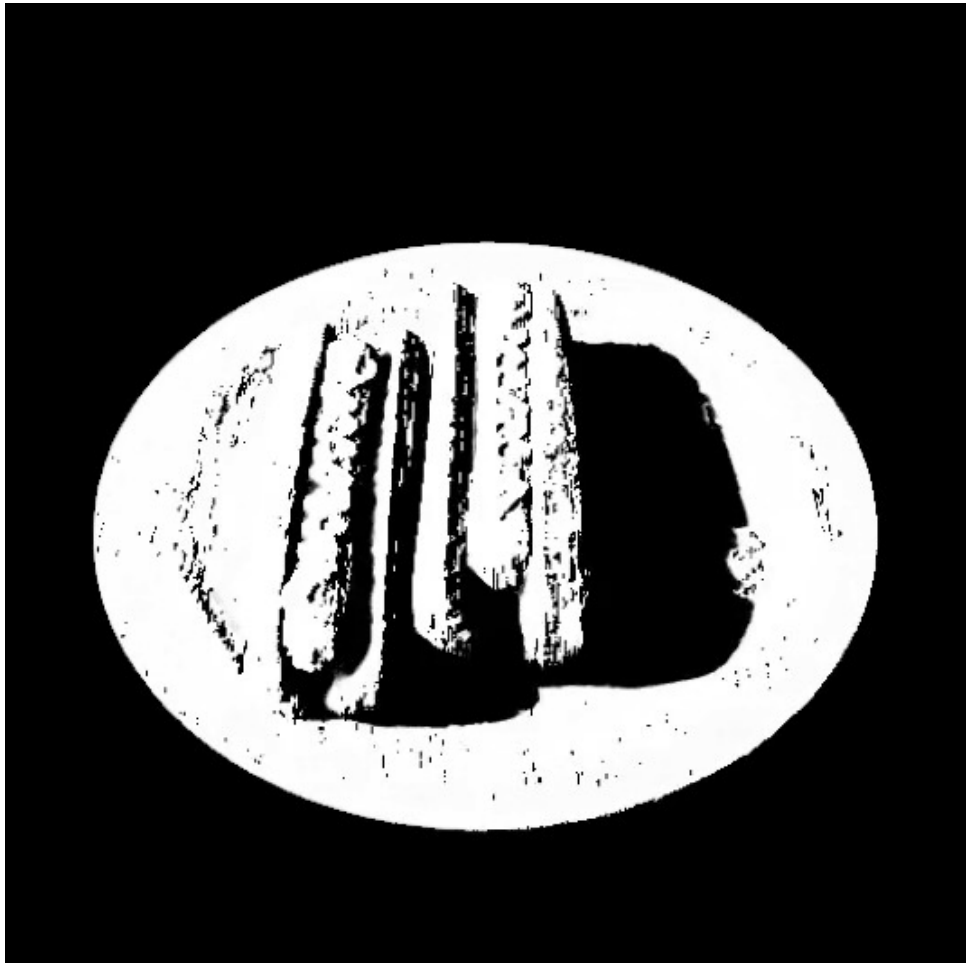

NeRF's Light Visibility

NeRFactor

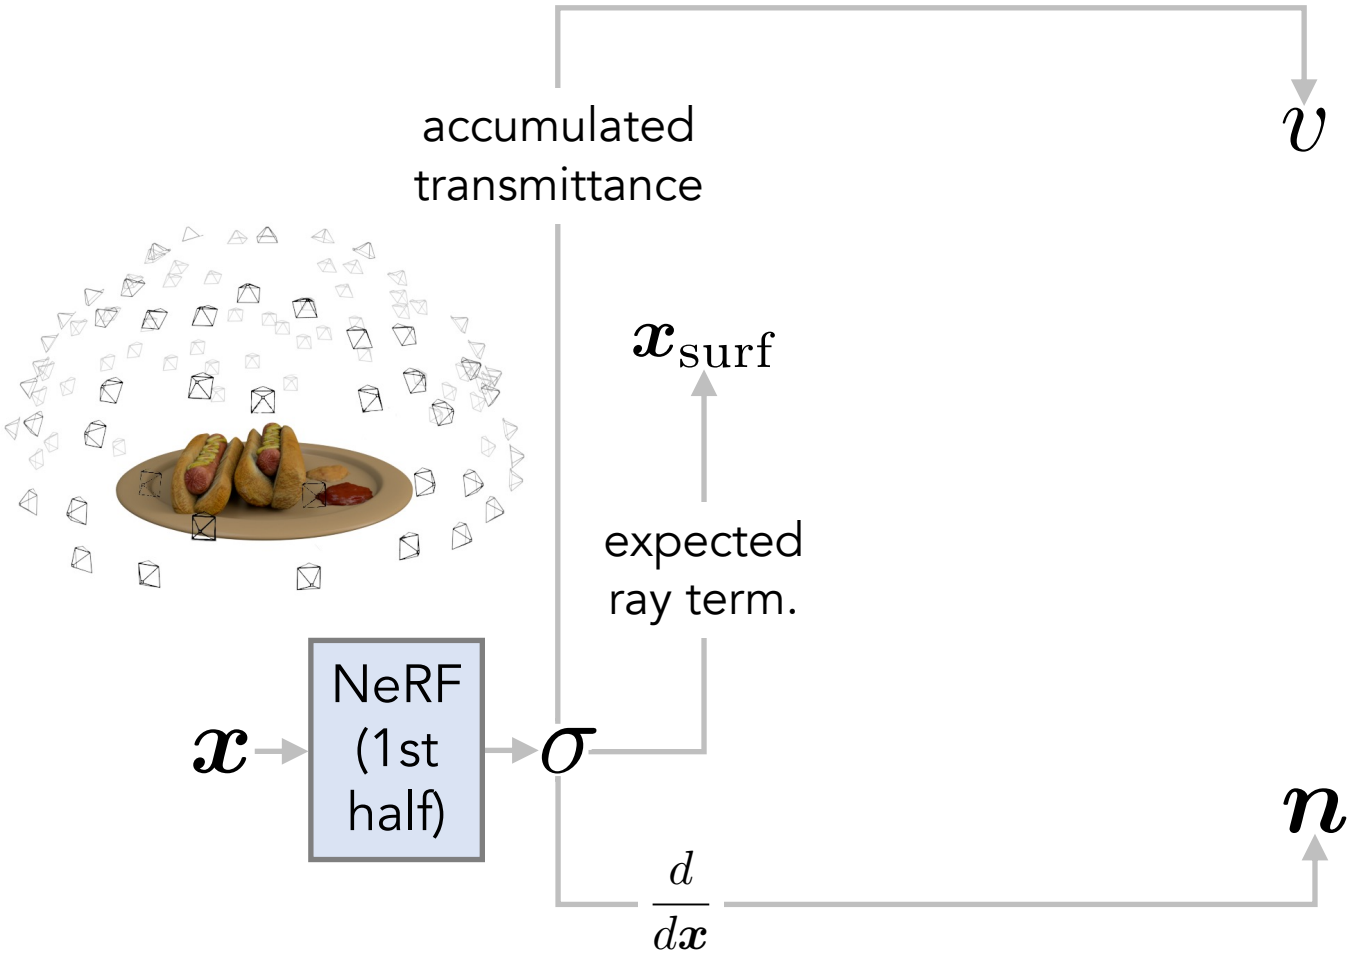
accumulated
transmittance
v
$\boldsymbol{x}_{\text{surf}}$
expected
ray term.
$\boldsymbol{x}$
NeRF
(1st
half)
σ
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$

NeRFactor

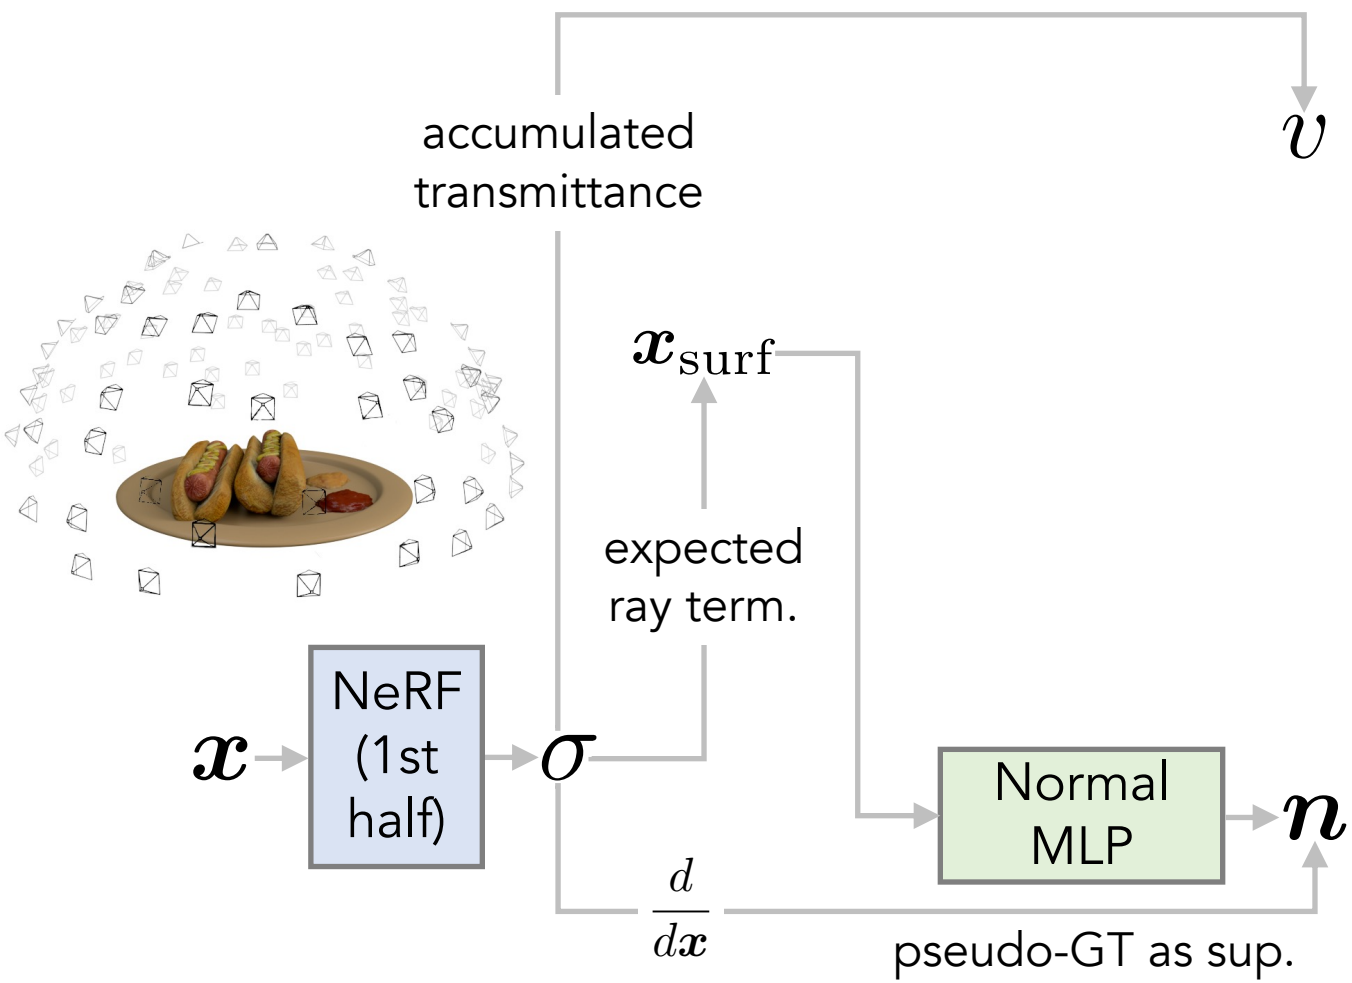
: pre-trained; frozen
: pre-trained; jointly finetuned
accumulated transmittance
v
x_{surf}
expected ray term.
x
NeRF (1st half)
σ
Normal MLP
n
$\frac{d}{dx}$
pseudo-GT as sup.

NeRFactor

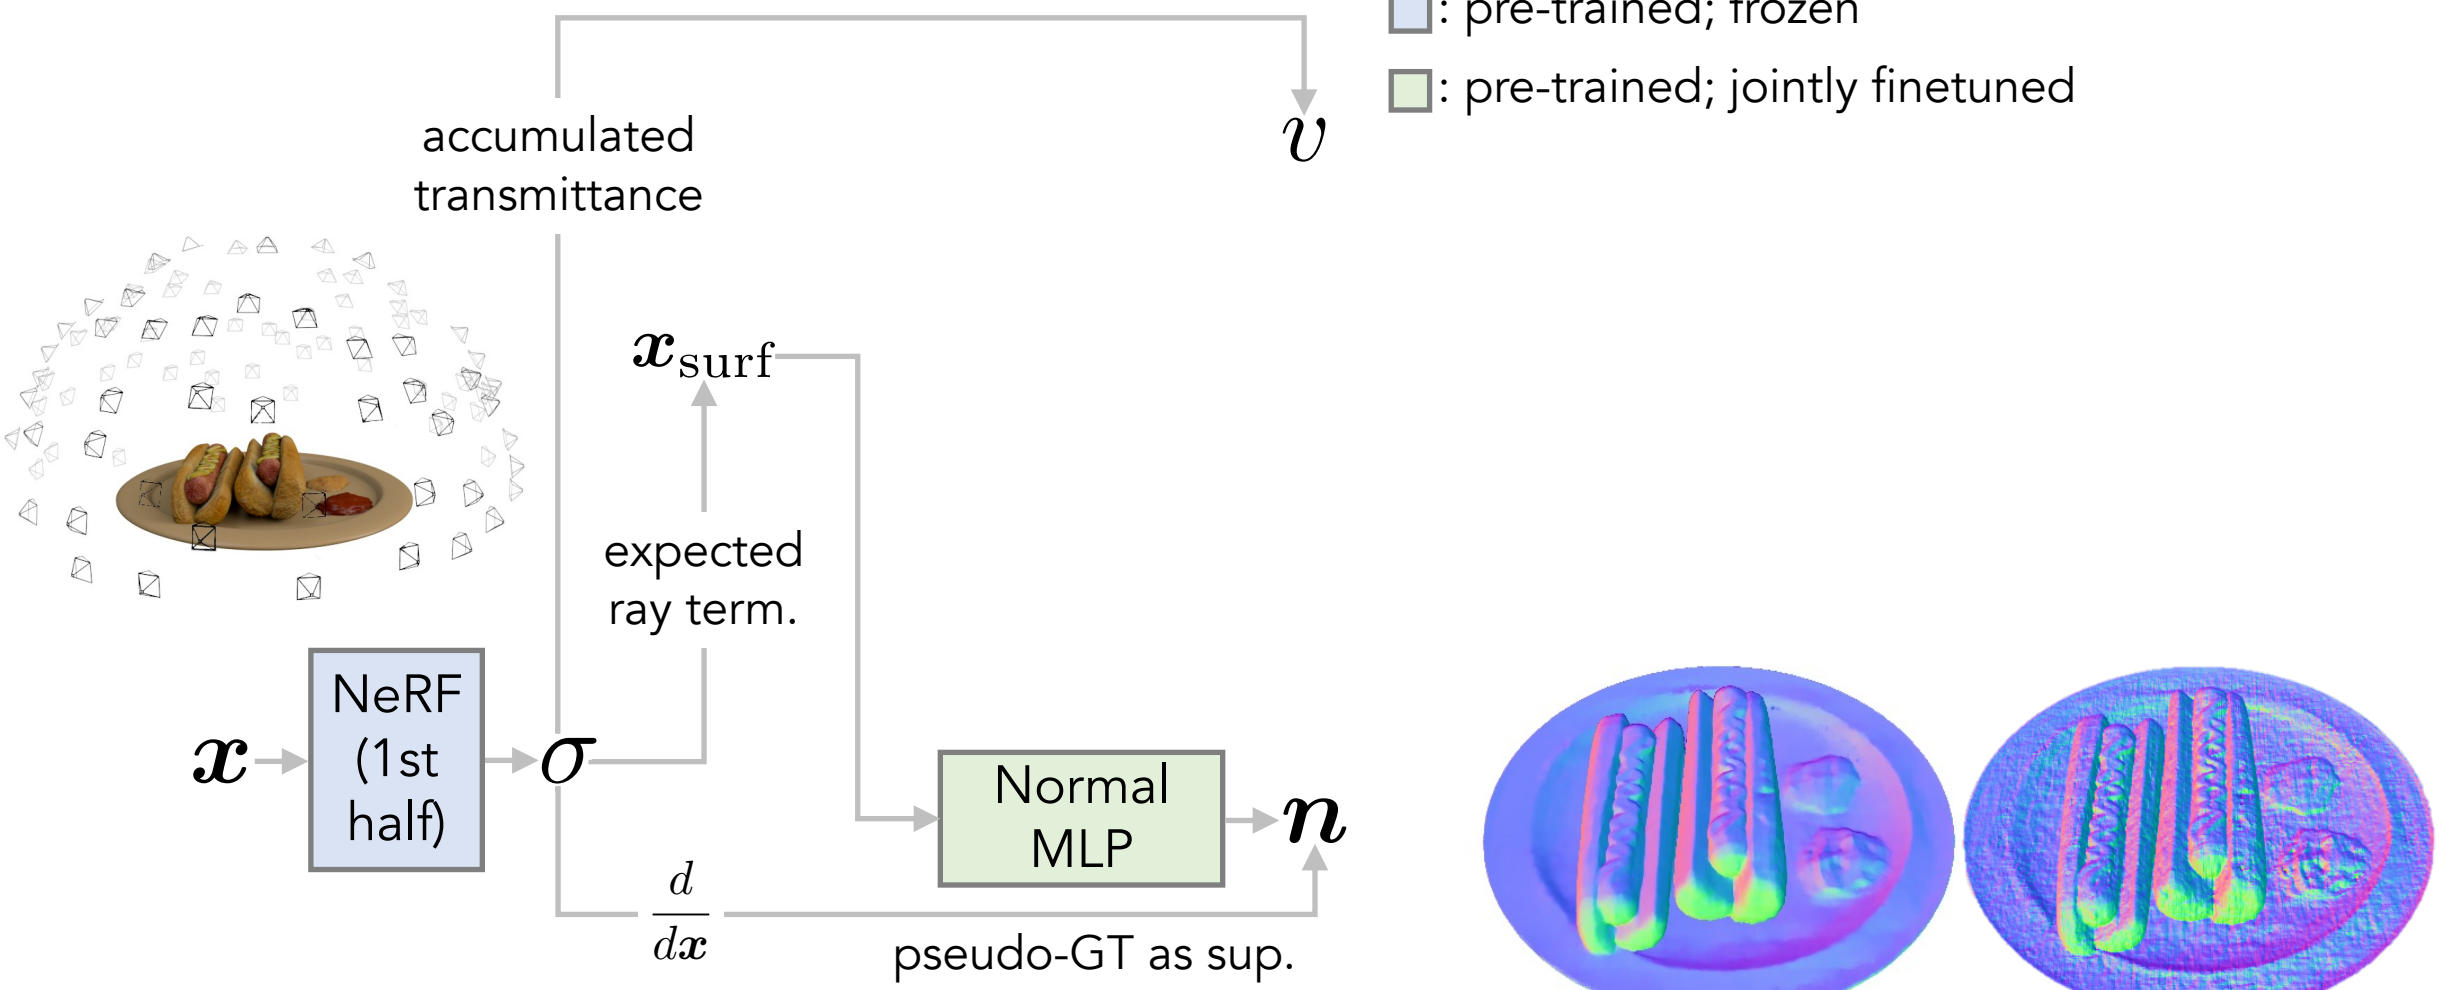
: pre-trained; frozen
: pre-trained; jointly finetuned
accumulated transmittance
v
x_{surf}
expected ray term.
x
NeRF (1st half)
σ
Normal MLP
n
$\frac{d}{dx}$
pseudo-GT as sup.
NeRFactor
NeRF

NeRFactor

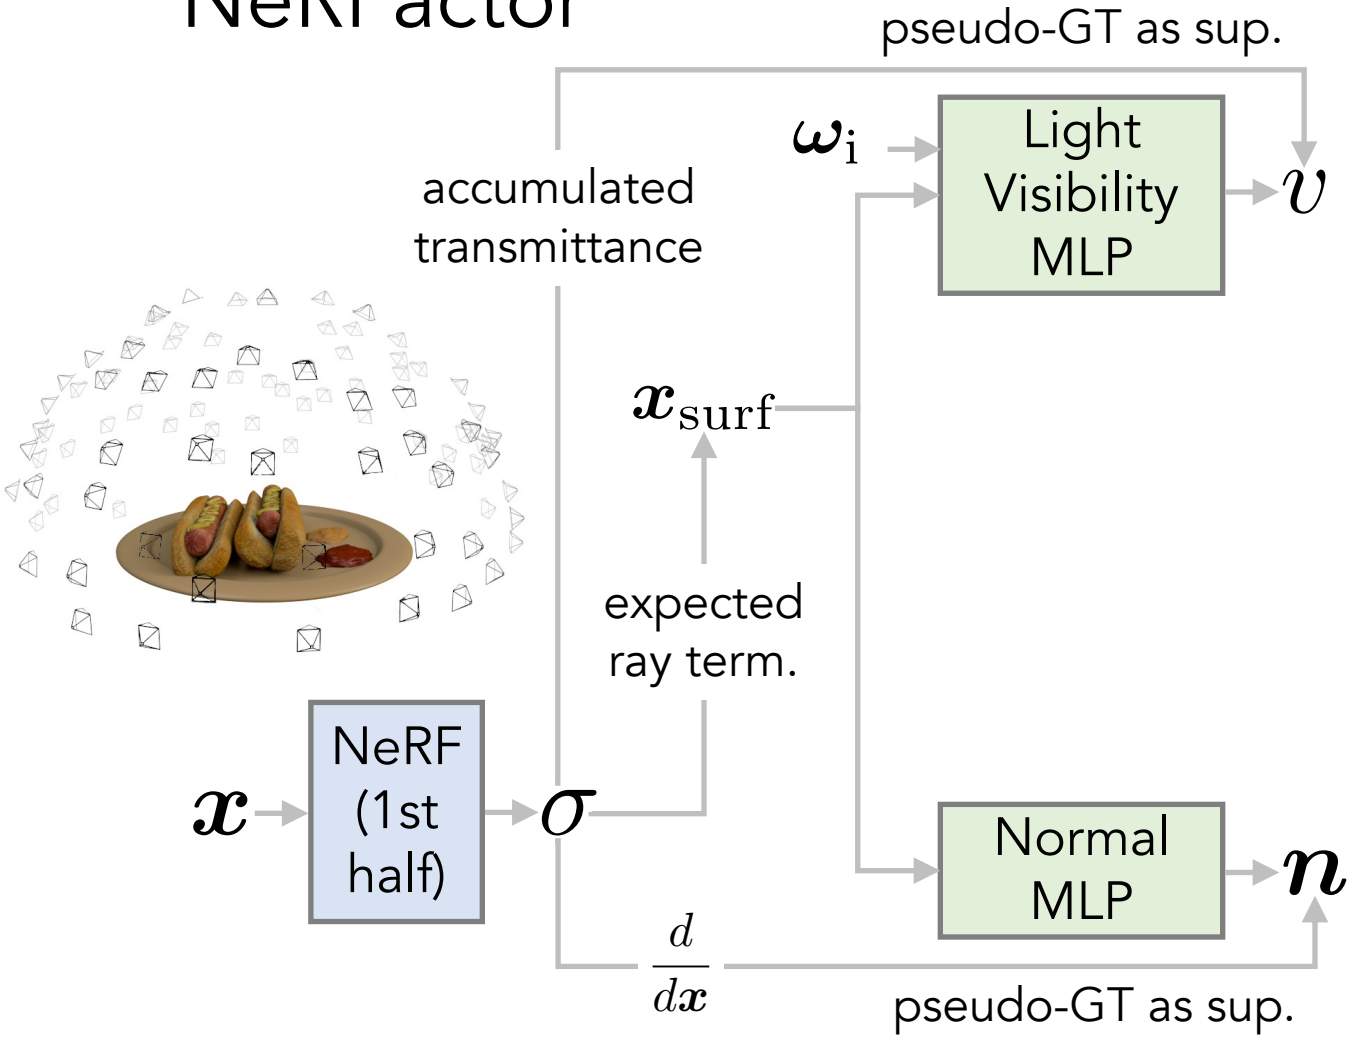

: pre-trained; frozen

: pre-trained; jointly finetuned

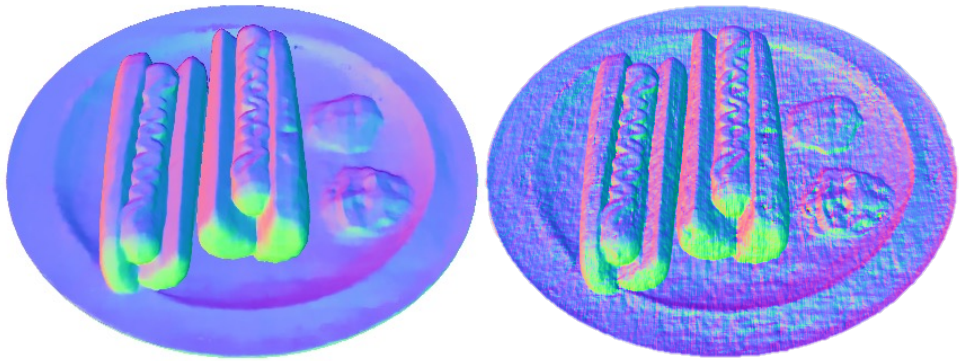

NeRFactor NeRF

NeRFactor

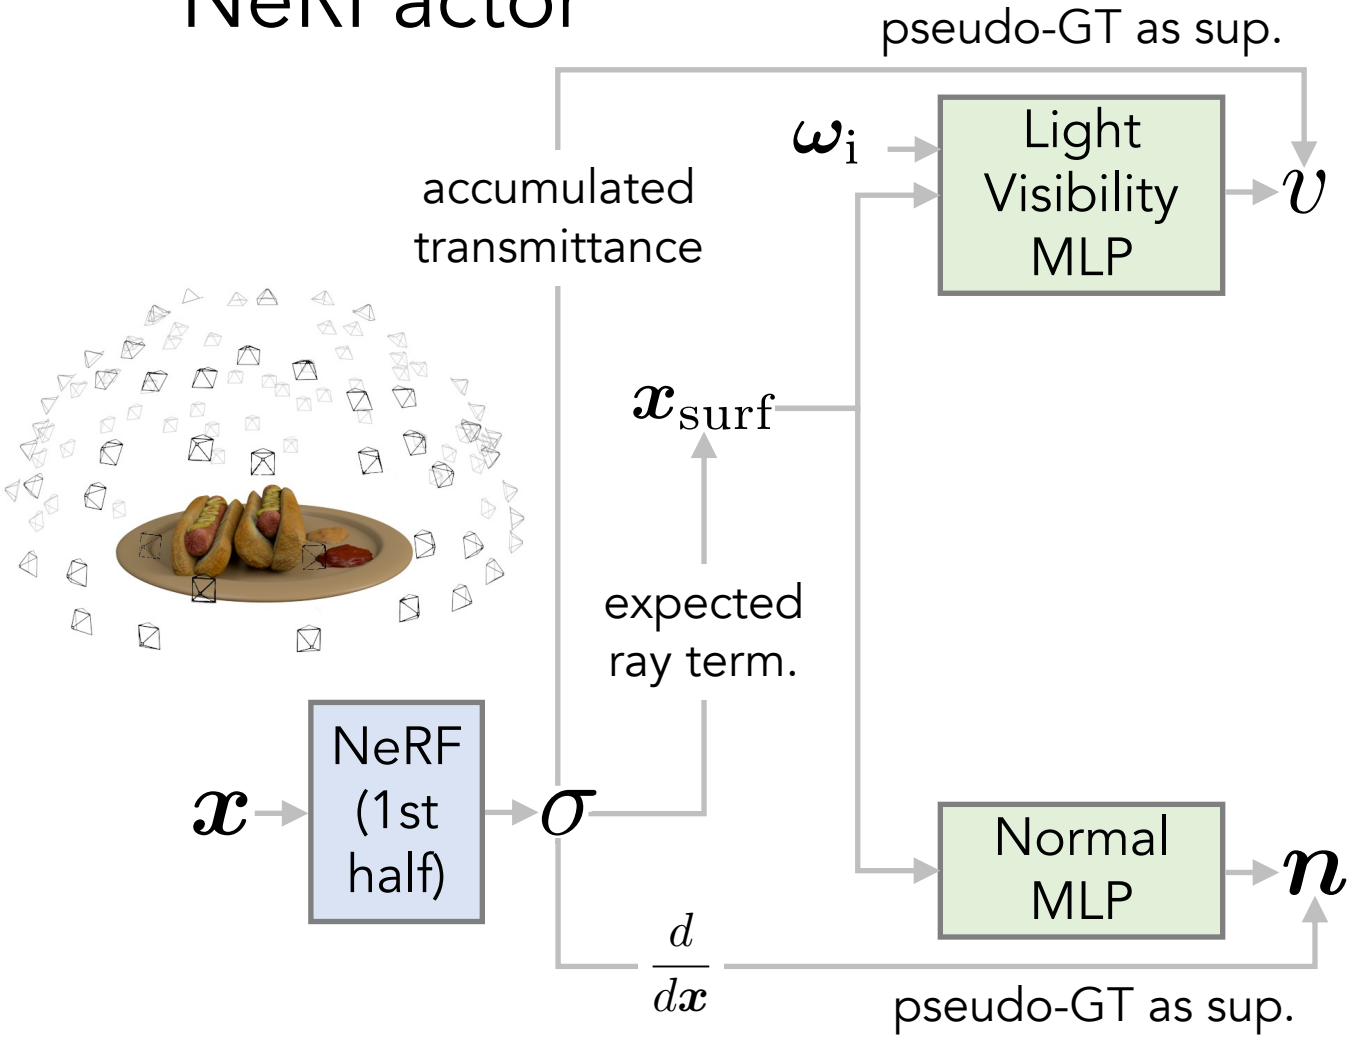

: pre-trained; frozen

: pre-trained; jointly finetuned

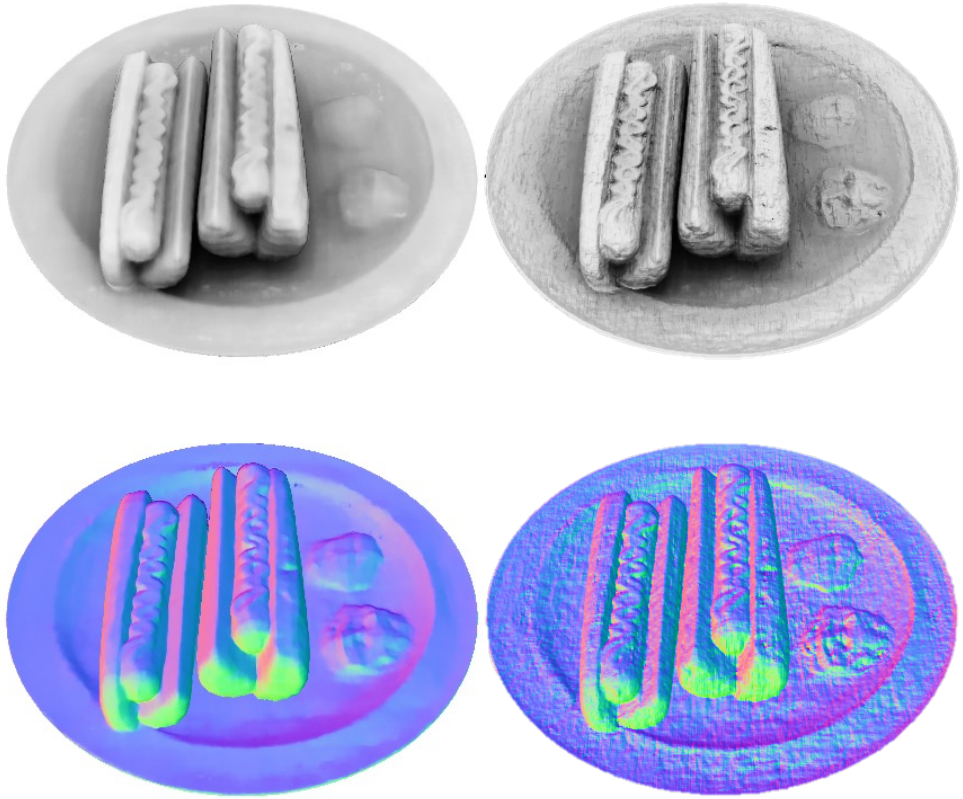

NeRFactor NeRF

NeRFactor

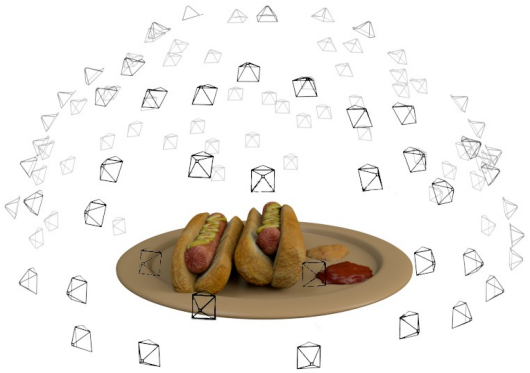

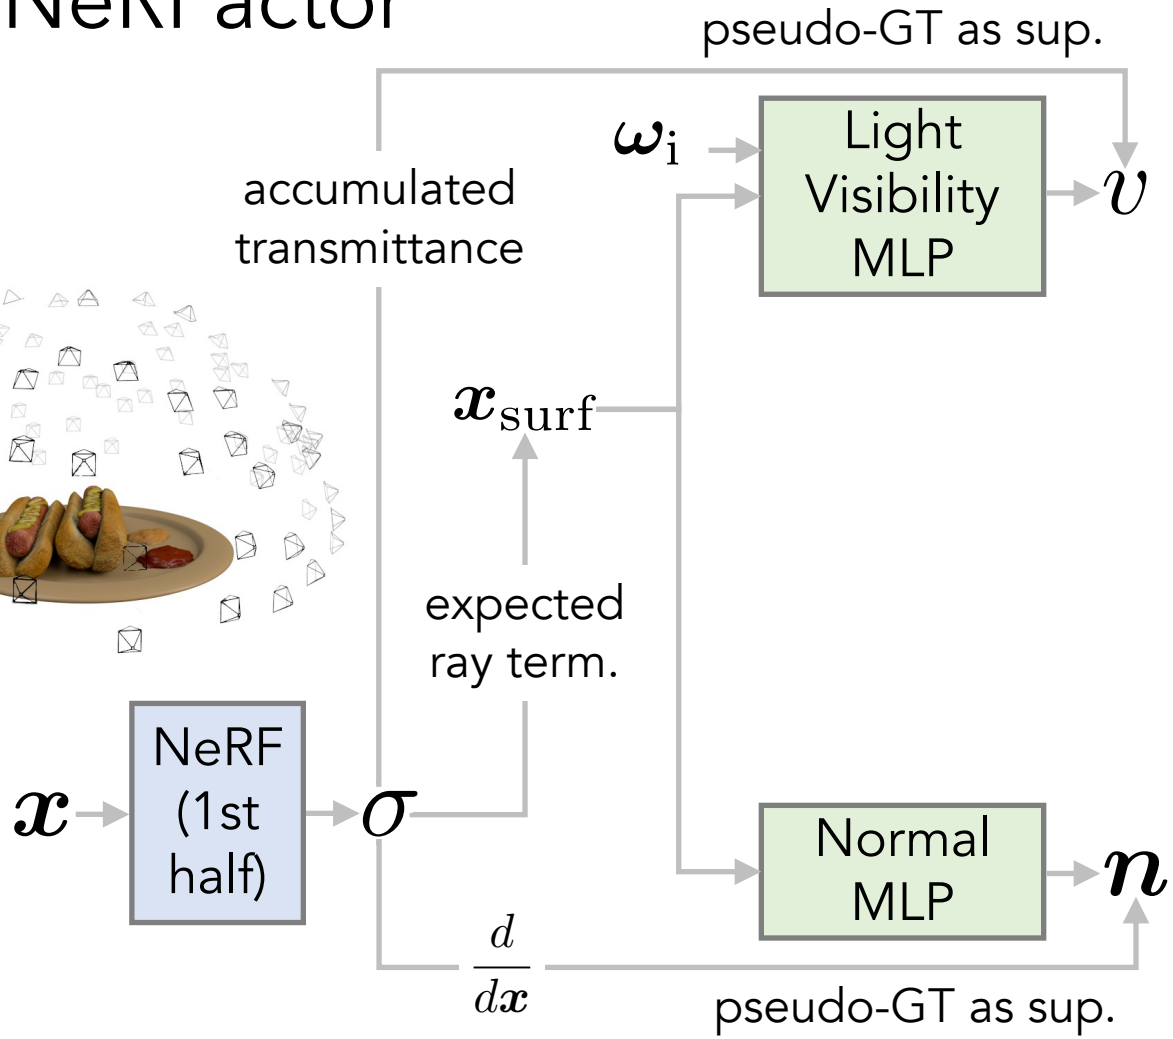

: pre-trained; frozen

: pre-trained; jointly finetuned

NeRFactor

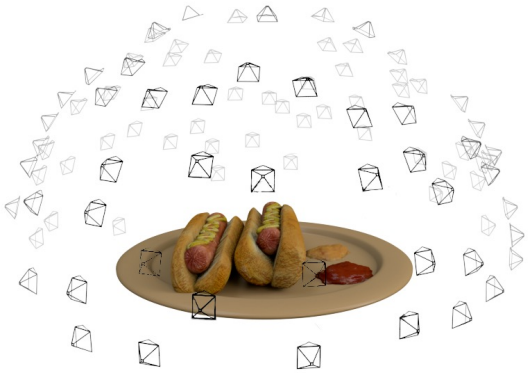

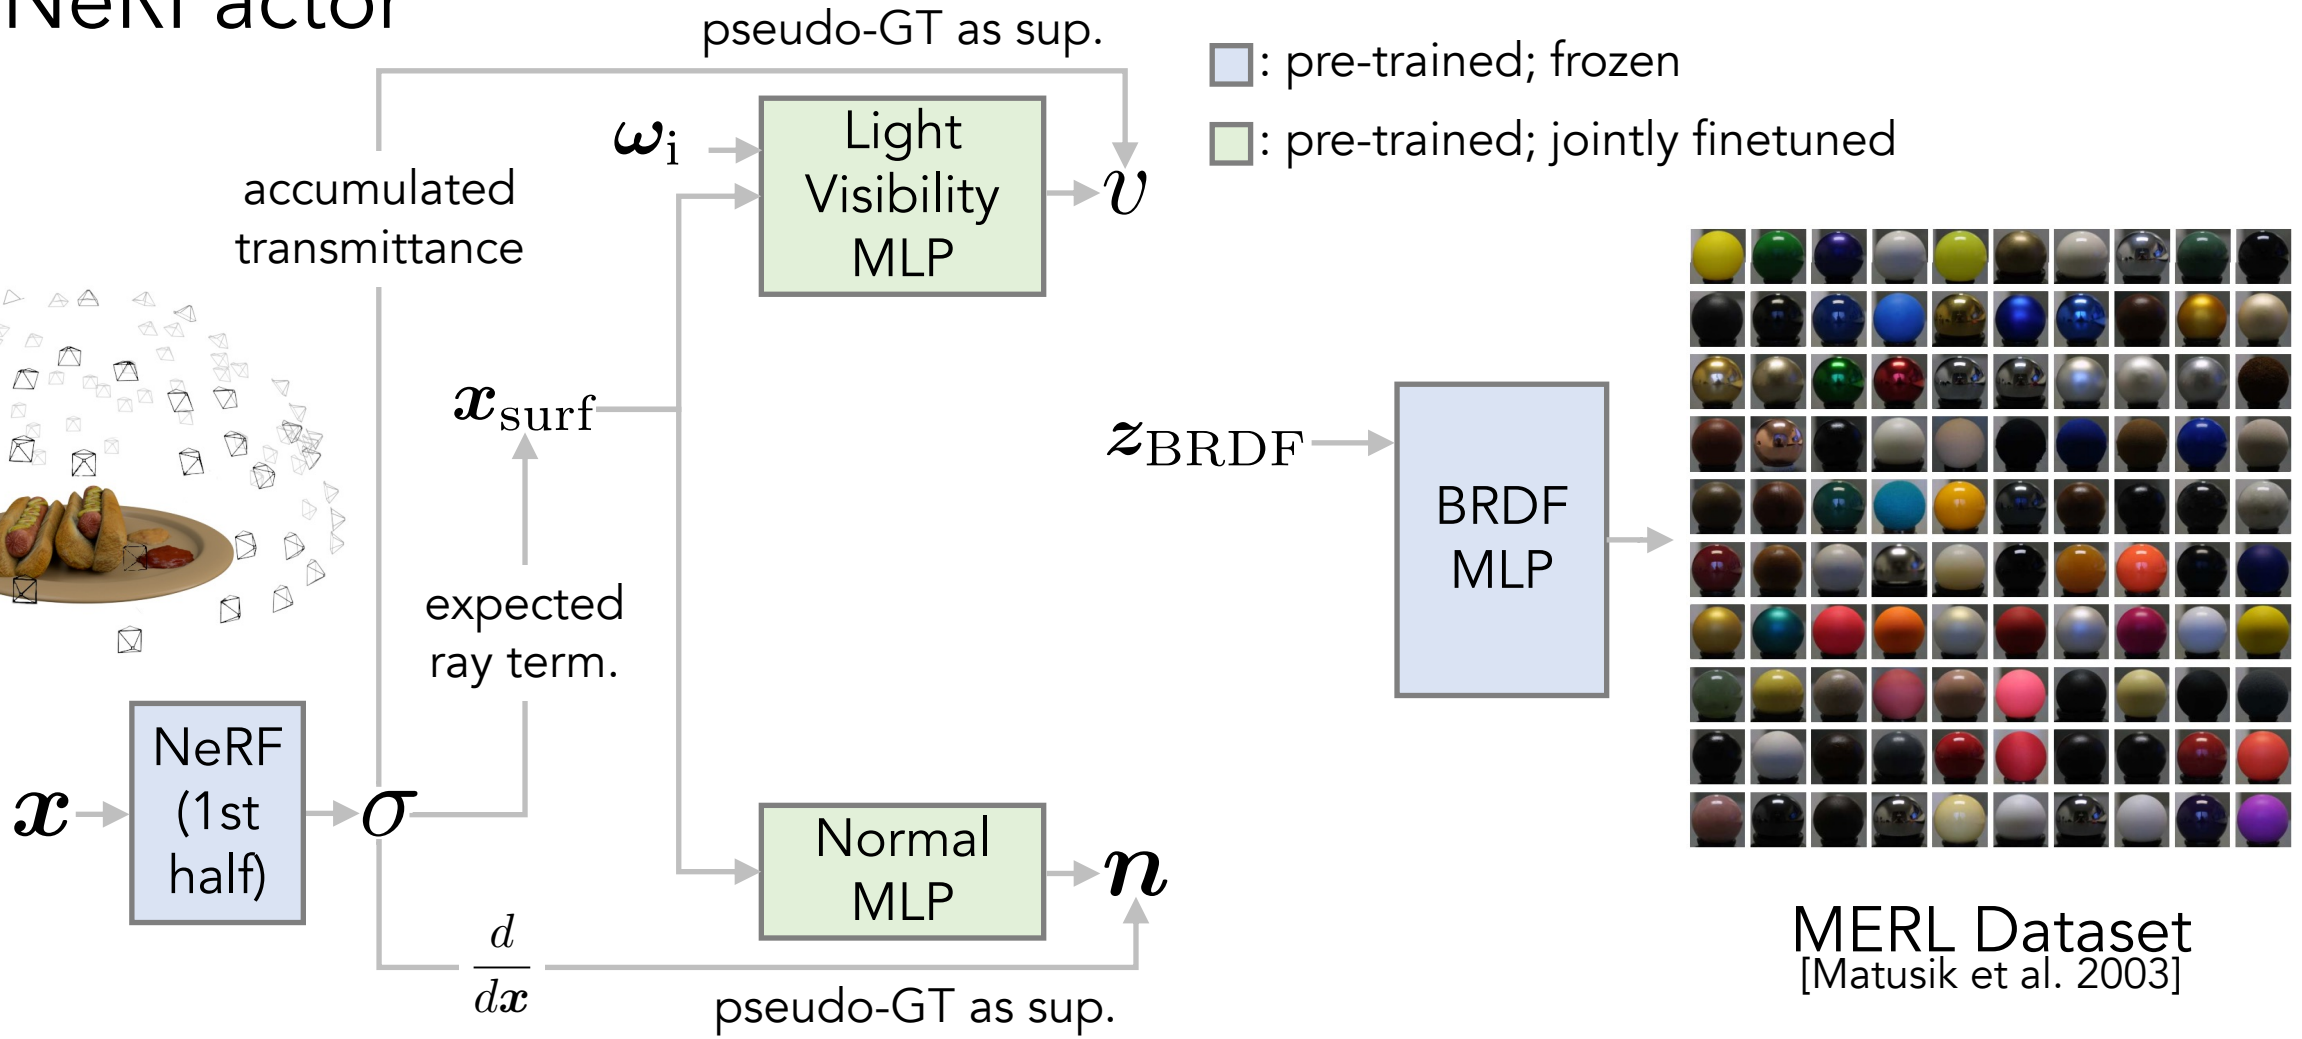

MERL Dataset
[Matusik et al. 2003]

Wojciech Matusik, Hanspeter Pfister, Matt Brand, Leonard McMillan
A Data-Driven Reflectance Model
TOG 2003

NeRFactor

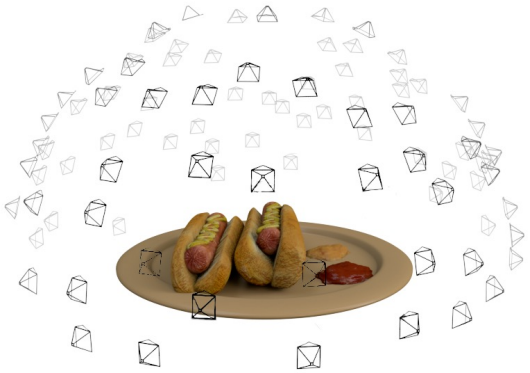

$\boldsymbol{x}$

NeRF (1st half)

σ

accumulated transmittance

expected ray term.

$\boldsymbol{x}_{\text{surf}}$

$\frac{d}{d\boldsymbol{x}}$

pseudo-GT as sup.

$\boldsymbol{\omega}_{\text{i}}$

Light Visibility MLP

$\mathcal{V}$

pseudo-GT as sup.

Normal MLP

$\boldsymbol{n}$

$\boldsymbol{z}_{\text{BRDF}}$

$(\boldsymbol{\omega}_{\text{i}}, \boldsymbol{\omega}_{\text{o}})$

$(\phi_{\text{d}}, \theta_{\text{h}}, \theta_{\text{d}})$

BRDF MLP

: pre-trained; frozen

: pre-trained; jointly finetuned

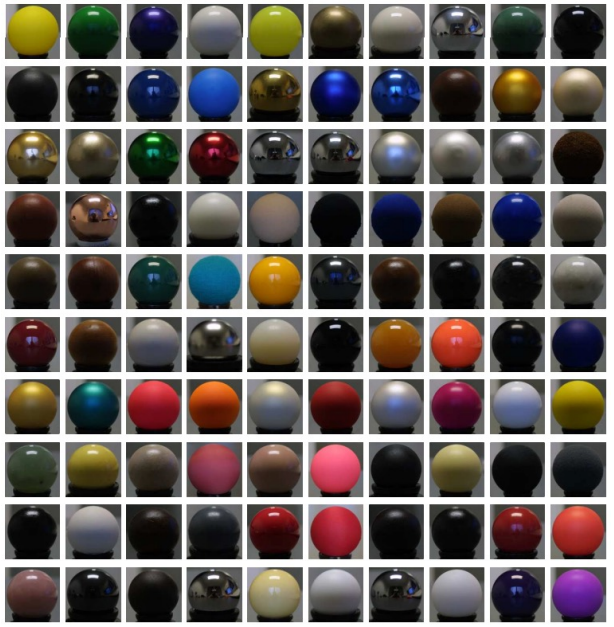

MERL Dataset
[Matusik et al. 2003]

Wojciech Matusik, Hanspeter Pfister, Matt Brand, Leonard McMillan
A Data-Driven Reflectance Model
TOG 2003

NeRFactor

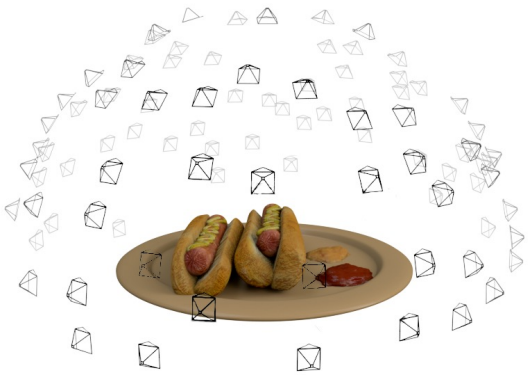

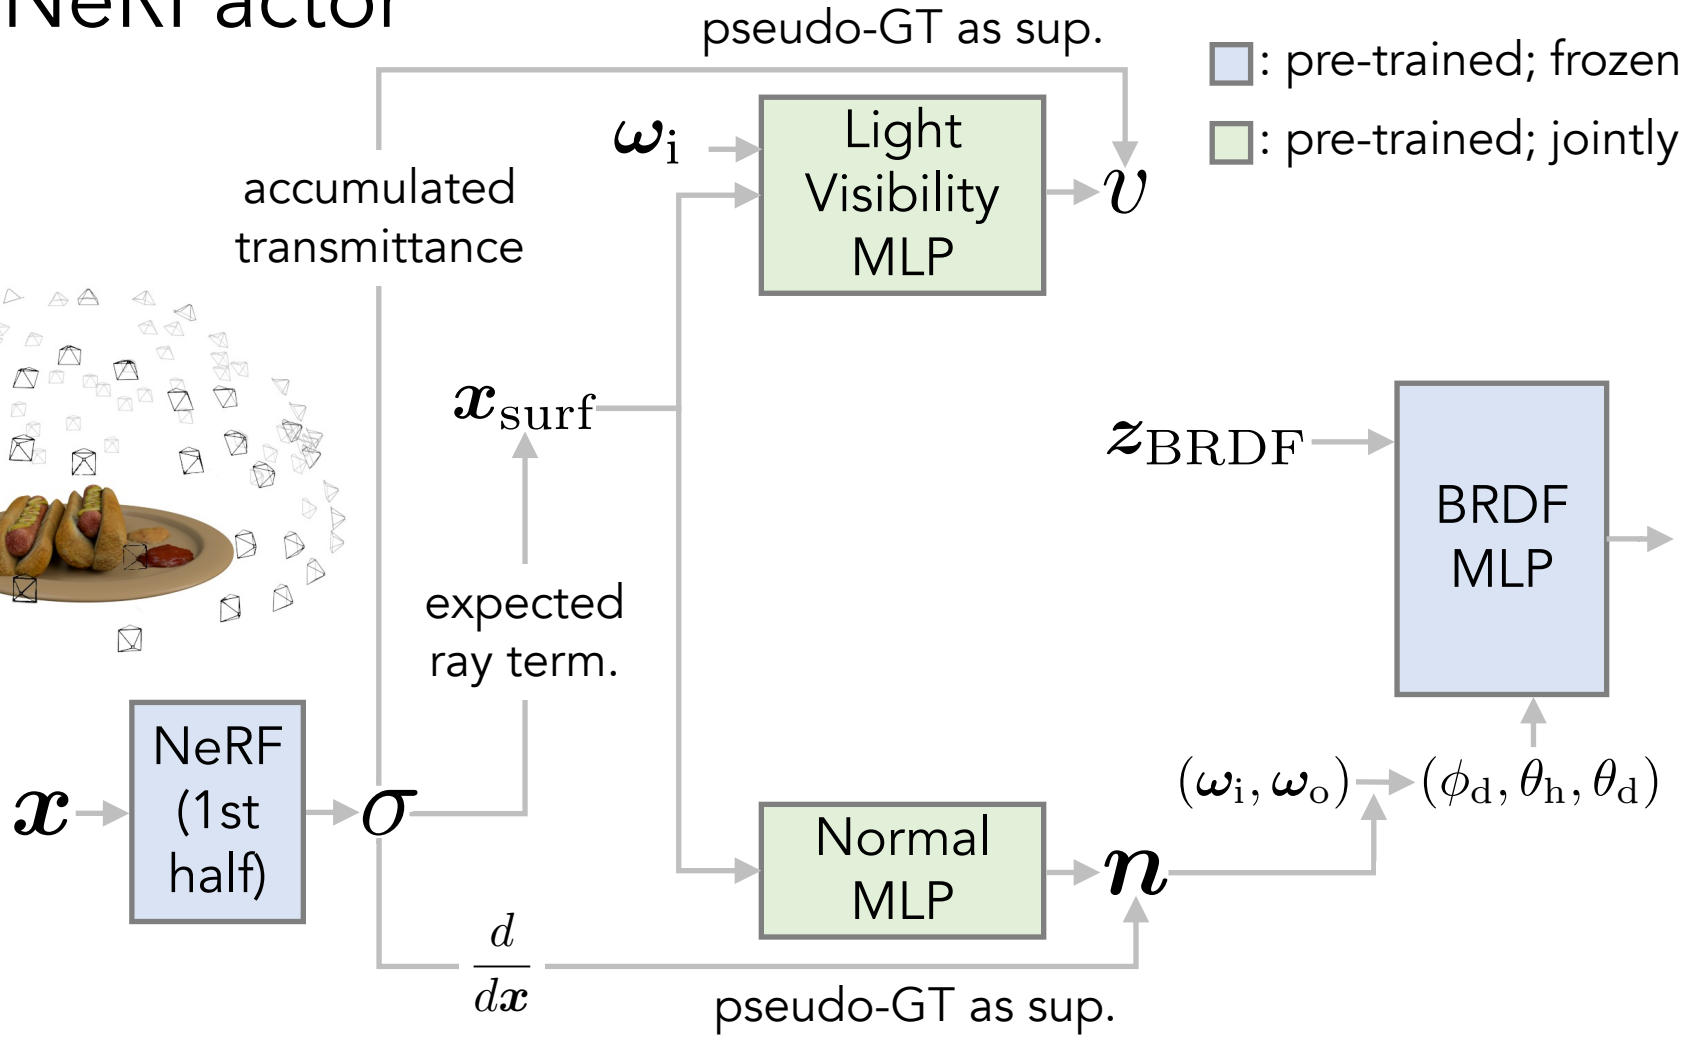
pseudo-GT as sup.
: pre-trained; frozen
: pre-trained; jointly finetuned
$\boldsymbol{\omega}_\mathrm{i}$
Light Visibility MLP
$\boldsymbol{v}$
accumulated transmittance
$\boldsymbol{x}_\mathrm{surf}$
$\boldsymbol{z}_\mathrm{BRDF}$
BRDF MLP
expected ray term.
$\boldsymbol{x}$
NeRF (1st half)
σ
$(\boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$
$(\phi_\mathrm{d}, \theta_\mathrm{h}, \theta_\mathrm{d})$
Normal MLP
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$
pseudo-GT as sup.

NeRFactor

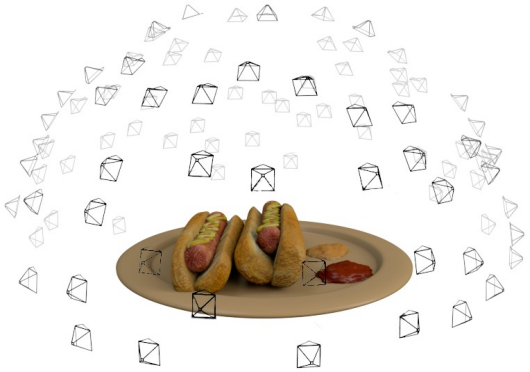

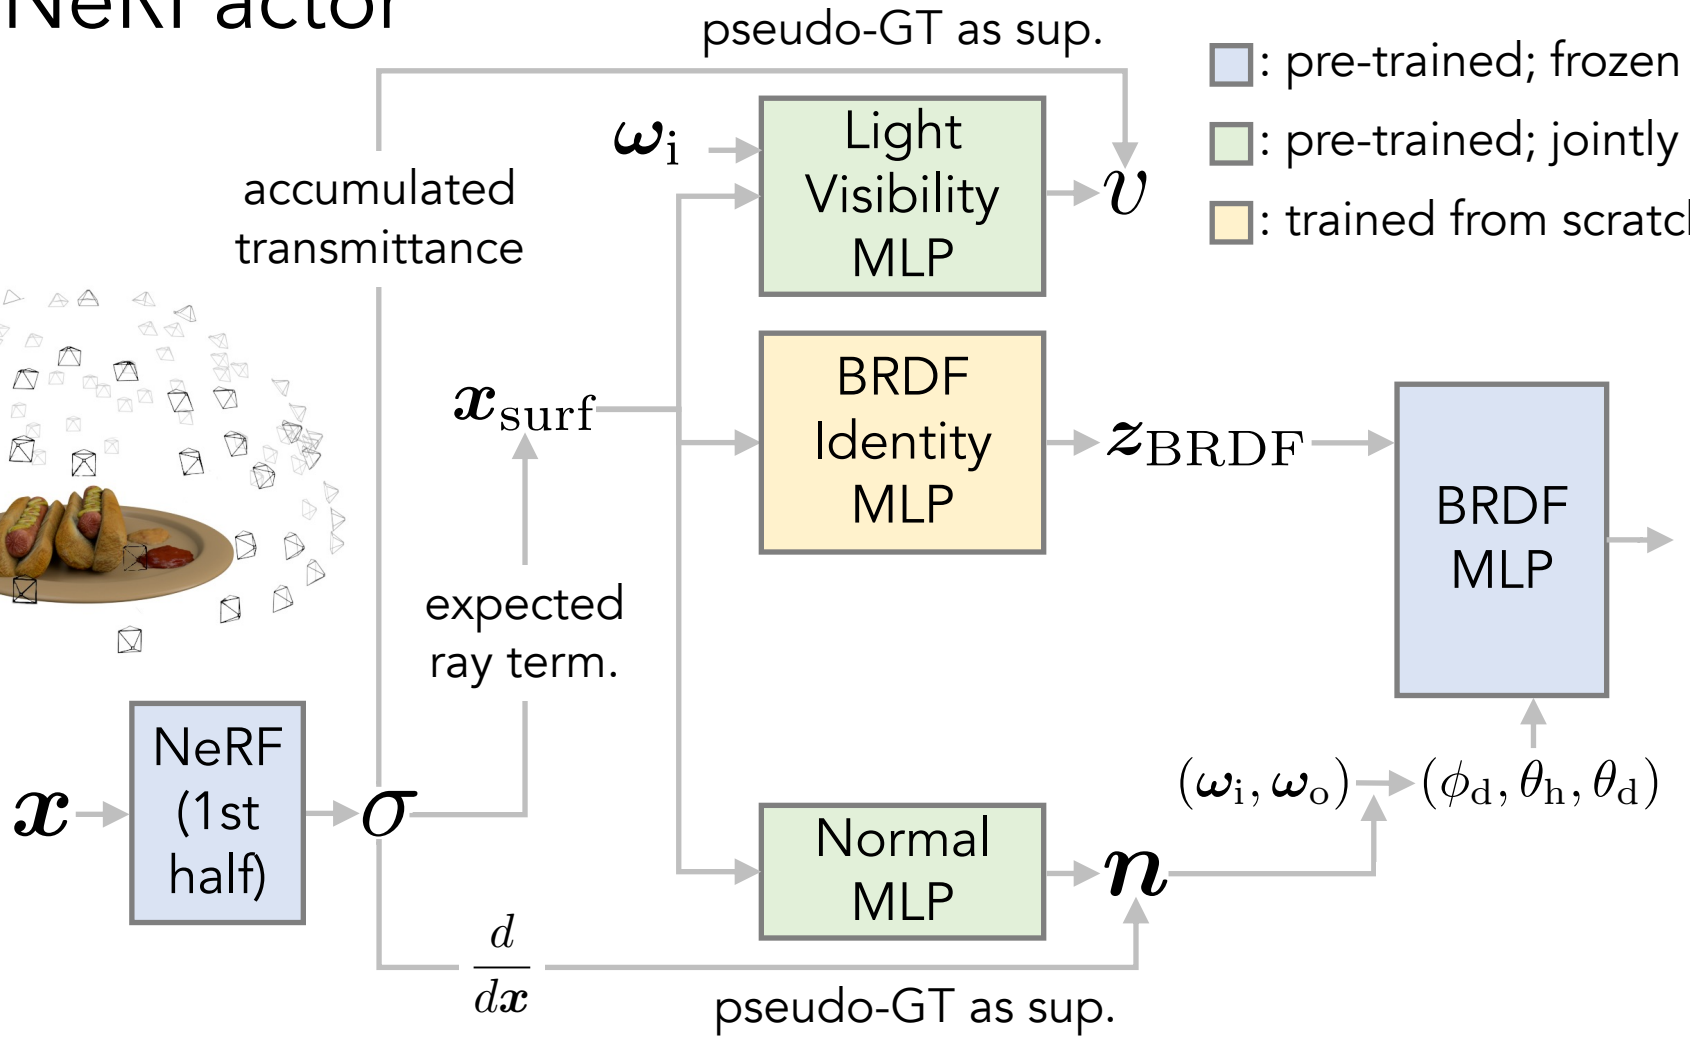
pseudo-GT as sup.
: pre-trained; frozen
: pre-trained; jointly finetuned
: trained from scratch
$\boldsymbol{\omega}_\mathrm{i}$
Light Visibility MLP
v
accumulated transmittance
$\boldsymbol{x}_\mathrm{surf}$
BRDF Identity MLP
$\boldsymbol{z}_\mathrm{BRDF}$
BRDF MLP
expected ray term.
$\boldsymbol{x}$
NeRF (1st half)
σ
$(\boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$
$(\phi_\mathrm{d}, \theta_\mathrm{h}, \theta_\mathrm{d})$
Normal MLP
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$
pseudo-GT as sup.

NeRFactor

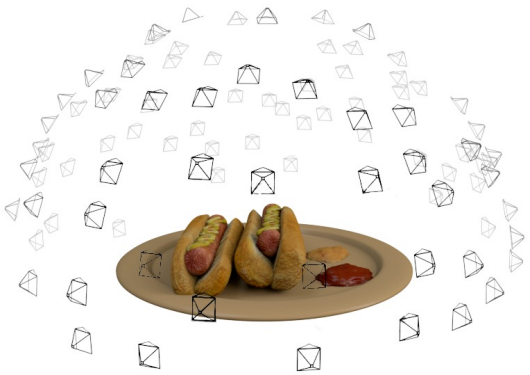

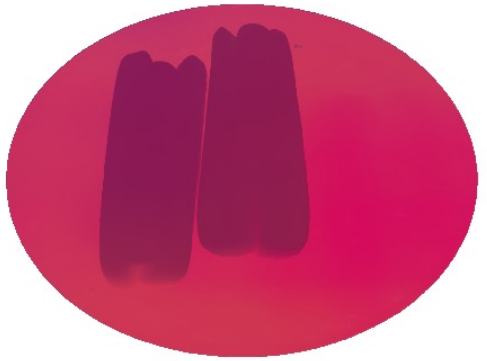

pseudo-GT as sup.
: pre-trained; frozen
: pre-trained; jointly finetuned
: trained from scratch
$\boldsymbol{\omega}_\mathrm{i}$
Light Visibility MLP
$\boldsymbol{v}$
accumulated transmittance
$\boldsymbol{x}_\mathrm{surf}$
BRDF Identity MLP
$\boldsymbol{z}_\mathrm{BRDF}$
BRDF MLP
$\boldsymbol{z}_\mathrm{BRDF}$
expected ray term.
$\boldsymbol{x}$
NeRF (1st half)
σ
$(\boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$
$(\phi_\mathrm{d}, \theta_\mathrm{h}, \theta_\mathrm{d})$
Normal MLP
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$
pseudo-GT as sup.

NeRFactor

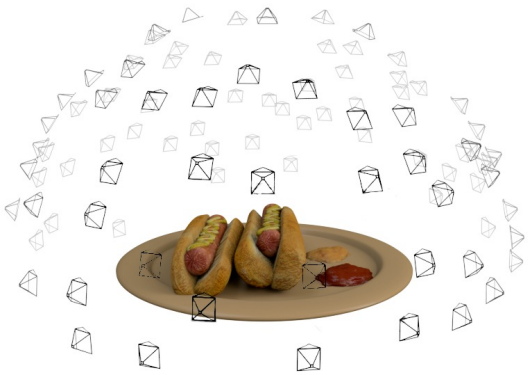

$\boldsymbol{x}$

NeRF (1st half)

σ

accumulated transmittance

expected ray term.

$\frac{d}{d\boldsymbol{x}}$

$\boldsymbol{x}_{\text{surf}}$

pseudo-GT as sup.

pseudo-GT as sup.

$\boldsymbol{\omega}_{\text{i}}$

Light Visibility MLP

v

BRDF Identity MLP

$\boldsymbol{z}_{\text{BRDF}}$

Albedo MLP

$\boldsymbol{a}$

Normal MLP

$\boldsymbol{n}$

$(\boldsymbol{\omega}_{\text{i}}, \boldsymbol{\omega}_{\text{o}})$

$(\phi_{\text{d}}, \theta_{\text{h}}, \theta_{\text{d}})$

BRDF MLP

: pre-trained; frozen

: pre-trained; jointly finetuned

: trained from scratch

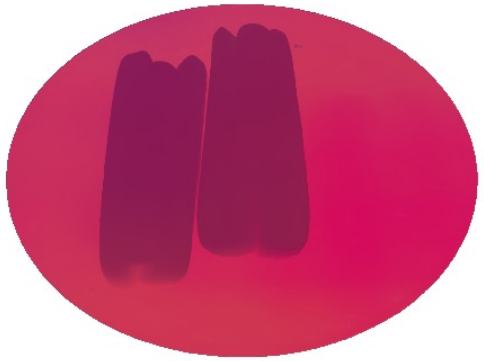

$\boldsymbol{z}_{\text{BRDF}}$

NeRFactor

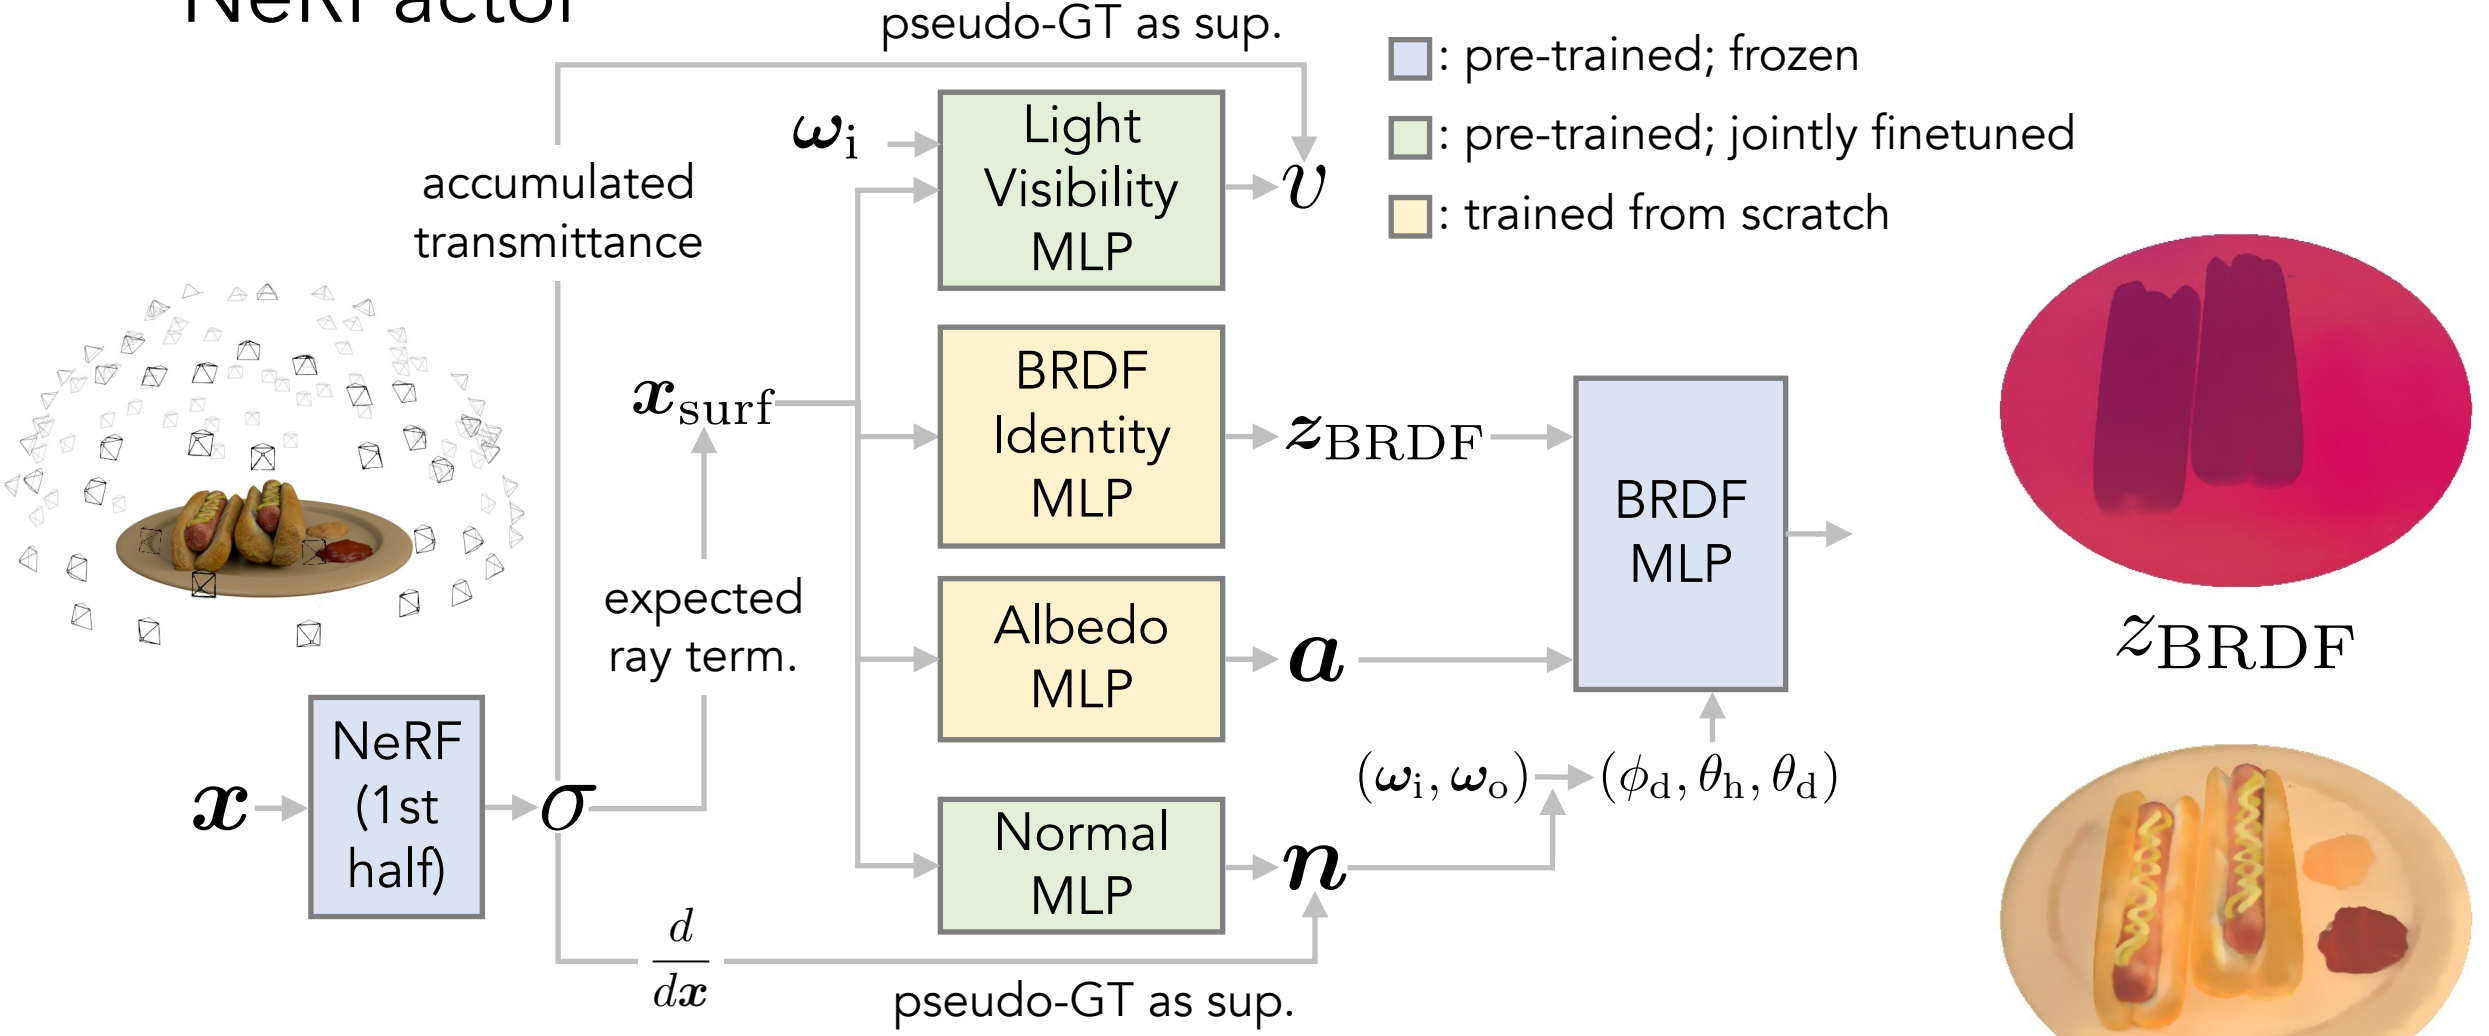

z_{BRDF}

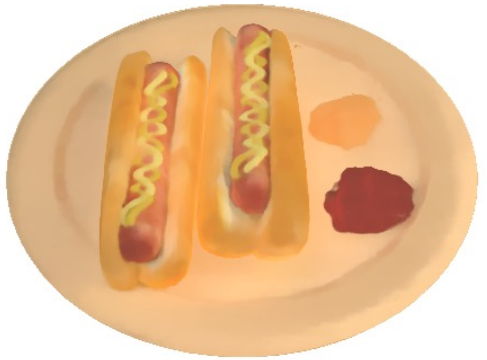

Albedo

NeRFactor

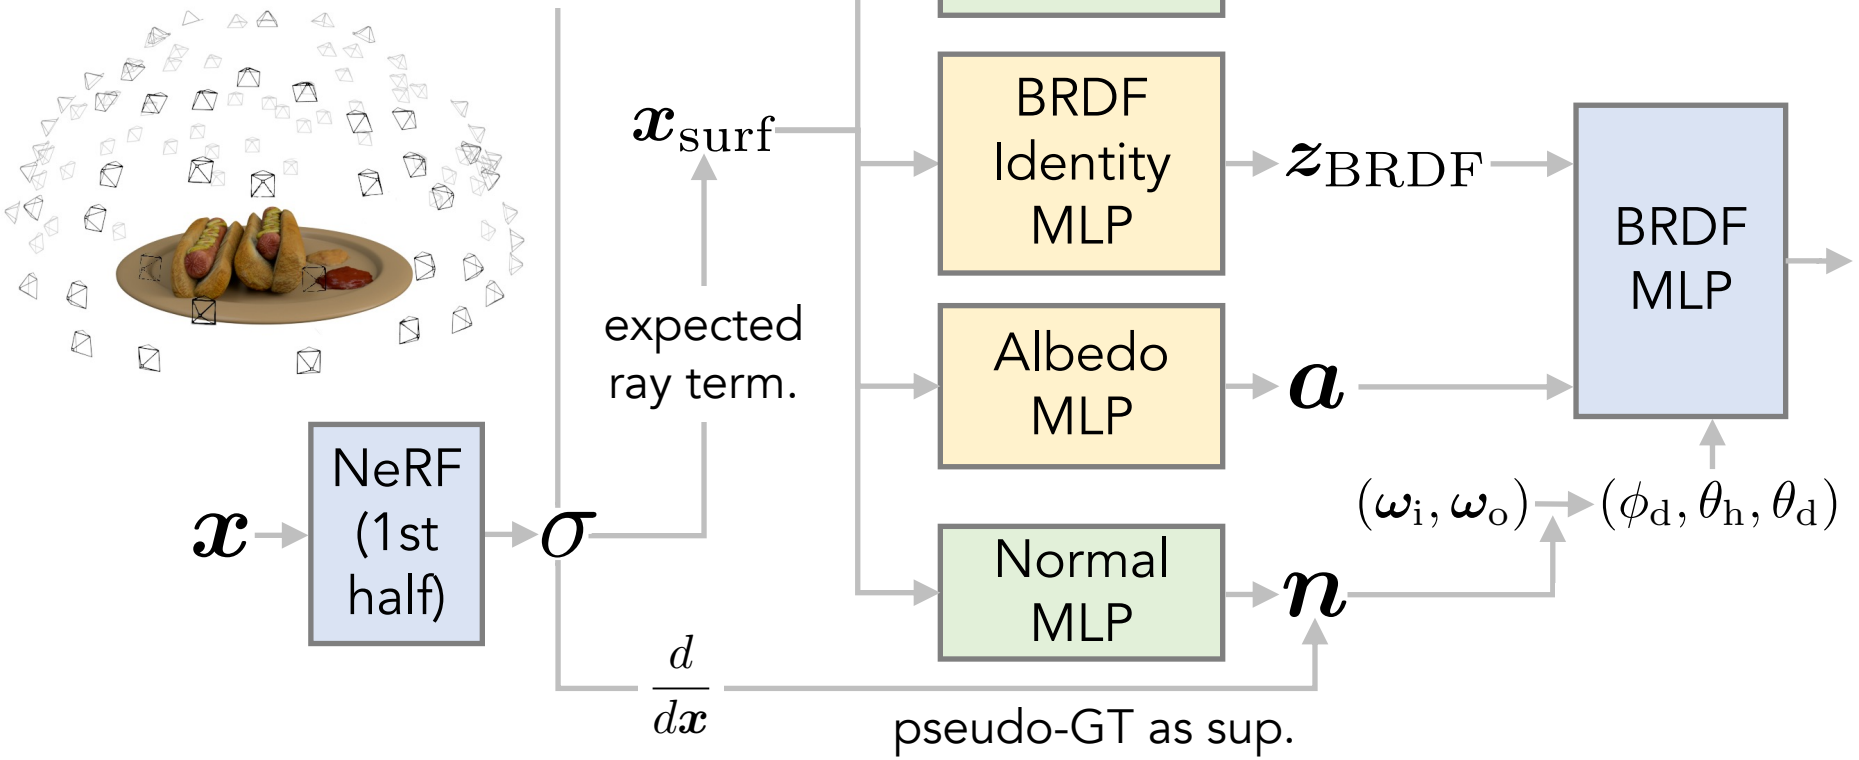
pseudo-GT as sup.
: pre-trained; frozen
: pre-trained; jointly finetuned
: trained from scratch
$\boldsymbol{\omega}_\mathrm{i}$
Light Visibility MLP
v
accumulated transmittance
$\boldsymbol{x}_\mathrm{surf}$
BRDF Identity MLP
$\boldsymbol{z}_\mathrm{BRDF}$
BRDF MLP
Albedo MLP
$\boldsymbol{a}$
expected ray term.
$\boldsymbol{x}$
NeRF (1st half)
σ
$(\boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$
$(\phi_\mathrm{d}, \theta_\mathrm{h}, \theta_\mathrm{d})$
Normal MLP
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$
pseudo-GT as sup.

NeRFactor

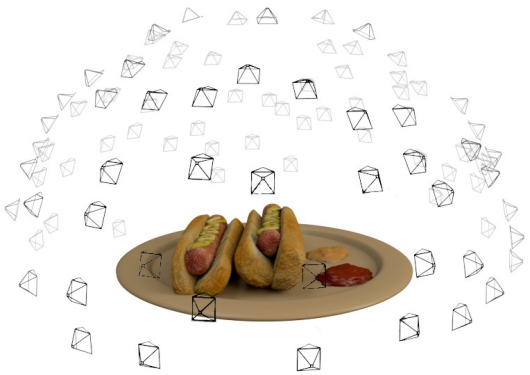

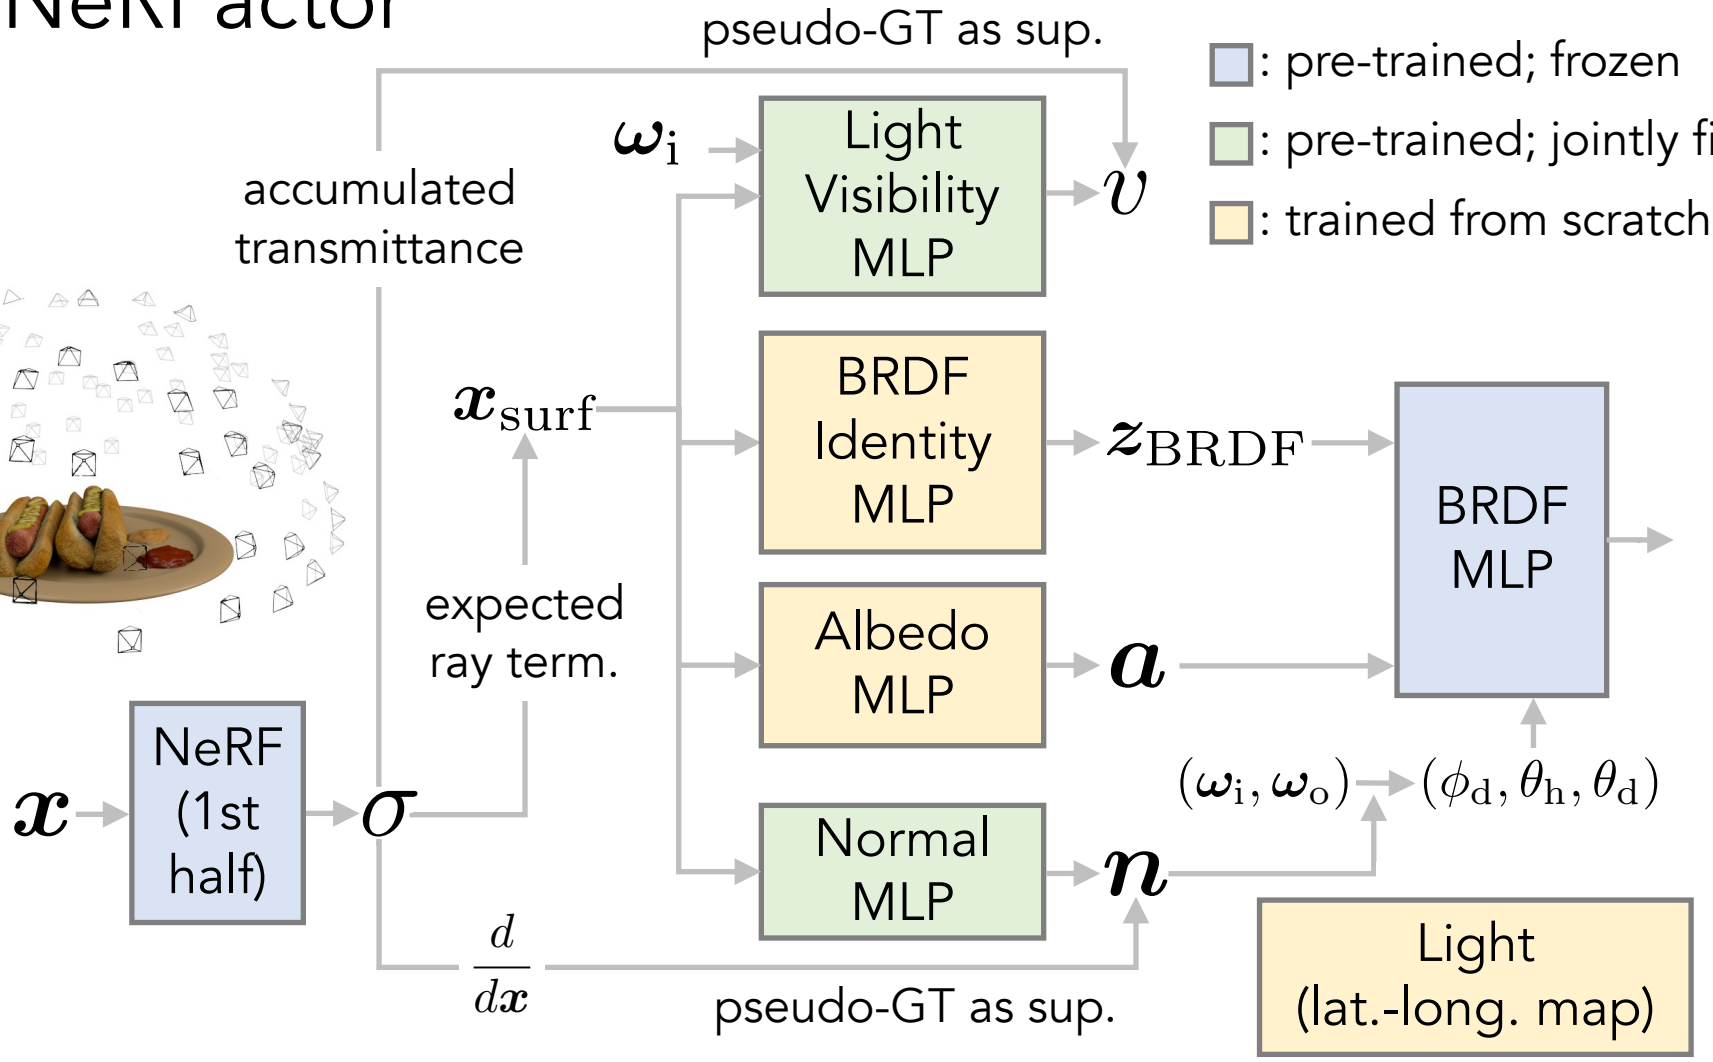
pseudo-GT as sup.
: pre-trained; frozen
: pre-trained; jointly finetuned
: trained from scratch
$\boldsymbol{\omega}_\mathrm{i}$
Light Visibility MLP
$\boldsymbol{v}$
accumulated transmittance
$\boldsymbol{x}_\mathrm{surf}$
BRDF Identity MLP
$\boldsymbol{z}_\mathrm{BRDF}$
BRDF MLP
Albedo MLP
$\boldsymbol{a}$
expected ray term.
$\boldsymbol{x}$
NeRF (1st half)
σ
$(\boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$
$(\phi_\mathrm{d}, \theta_\mathrm{h}, \theta_\mathrm{d})$
Normal MLP
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$
pseudo-GT as sup.
Light (lat.-long. map)

NeRFactor

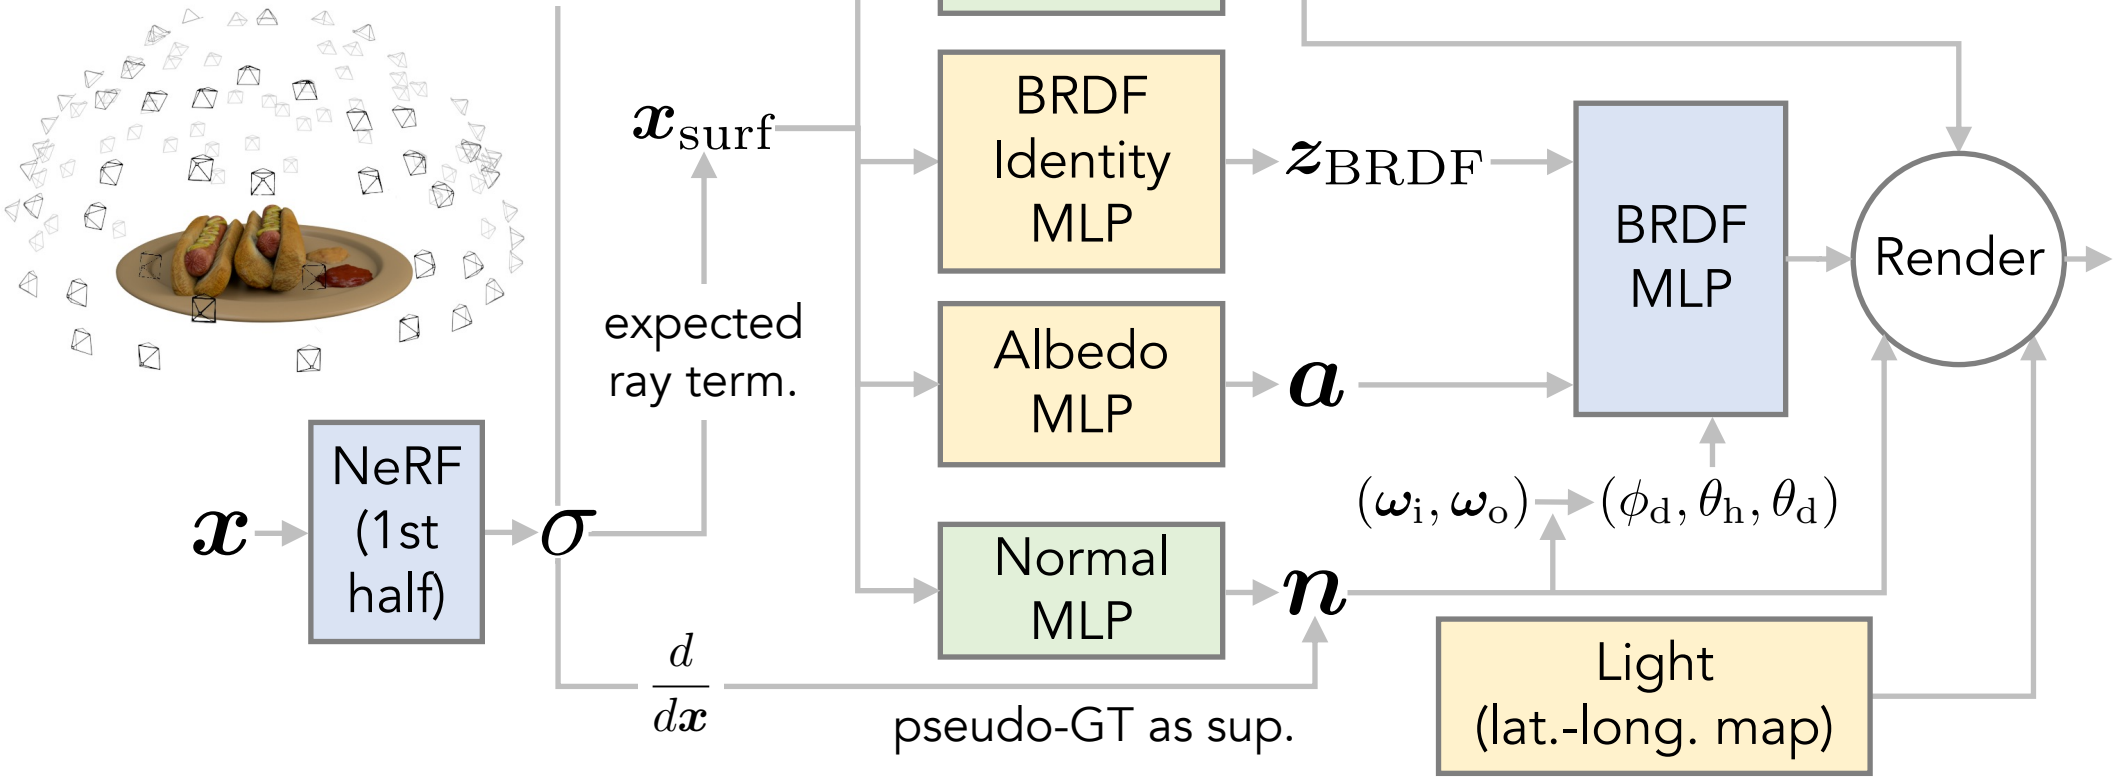
pseudo-GT as sup.
: pre-trained; frozen
: pre-trained; jointly finetuned
: trained from scratch
$\boldsymbol{\omega}_\mathrm{i}$
Light Visibility MLP
$\boldsymbol{v}$
accumulated transmittance
$\boldsymbol{x}_\mathrm{surf}$
BRDF Identity MLP
$\boldsymbol{z}_\mathrm{BRDF}$
BRDF MLP
Render
Albedo MLP
$\boldsymbol{a}$
expected ray term.
$\boldsymbol{x}$
NeRF (1st half)
σ
$(\boldsymbol{\omega}_\mathrm{i}, \boldsymbol{\omega}_\mathrm{o})$
$(\phi_\mathrm{d}, \theta_\mathrm{h}, \theta_\mathrm{d})$
Normal MLP
$\boldsymbol{n}$
$\frac{d}{d\boldsymbol{x}}$
pseudo-GT as sup.
Light (lat.-long. map)

NeRFactor

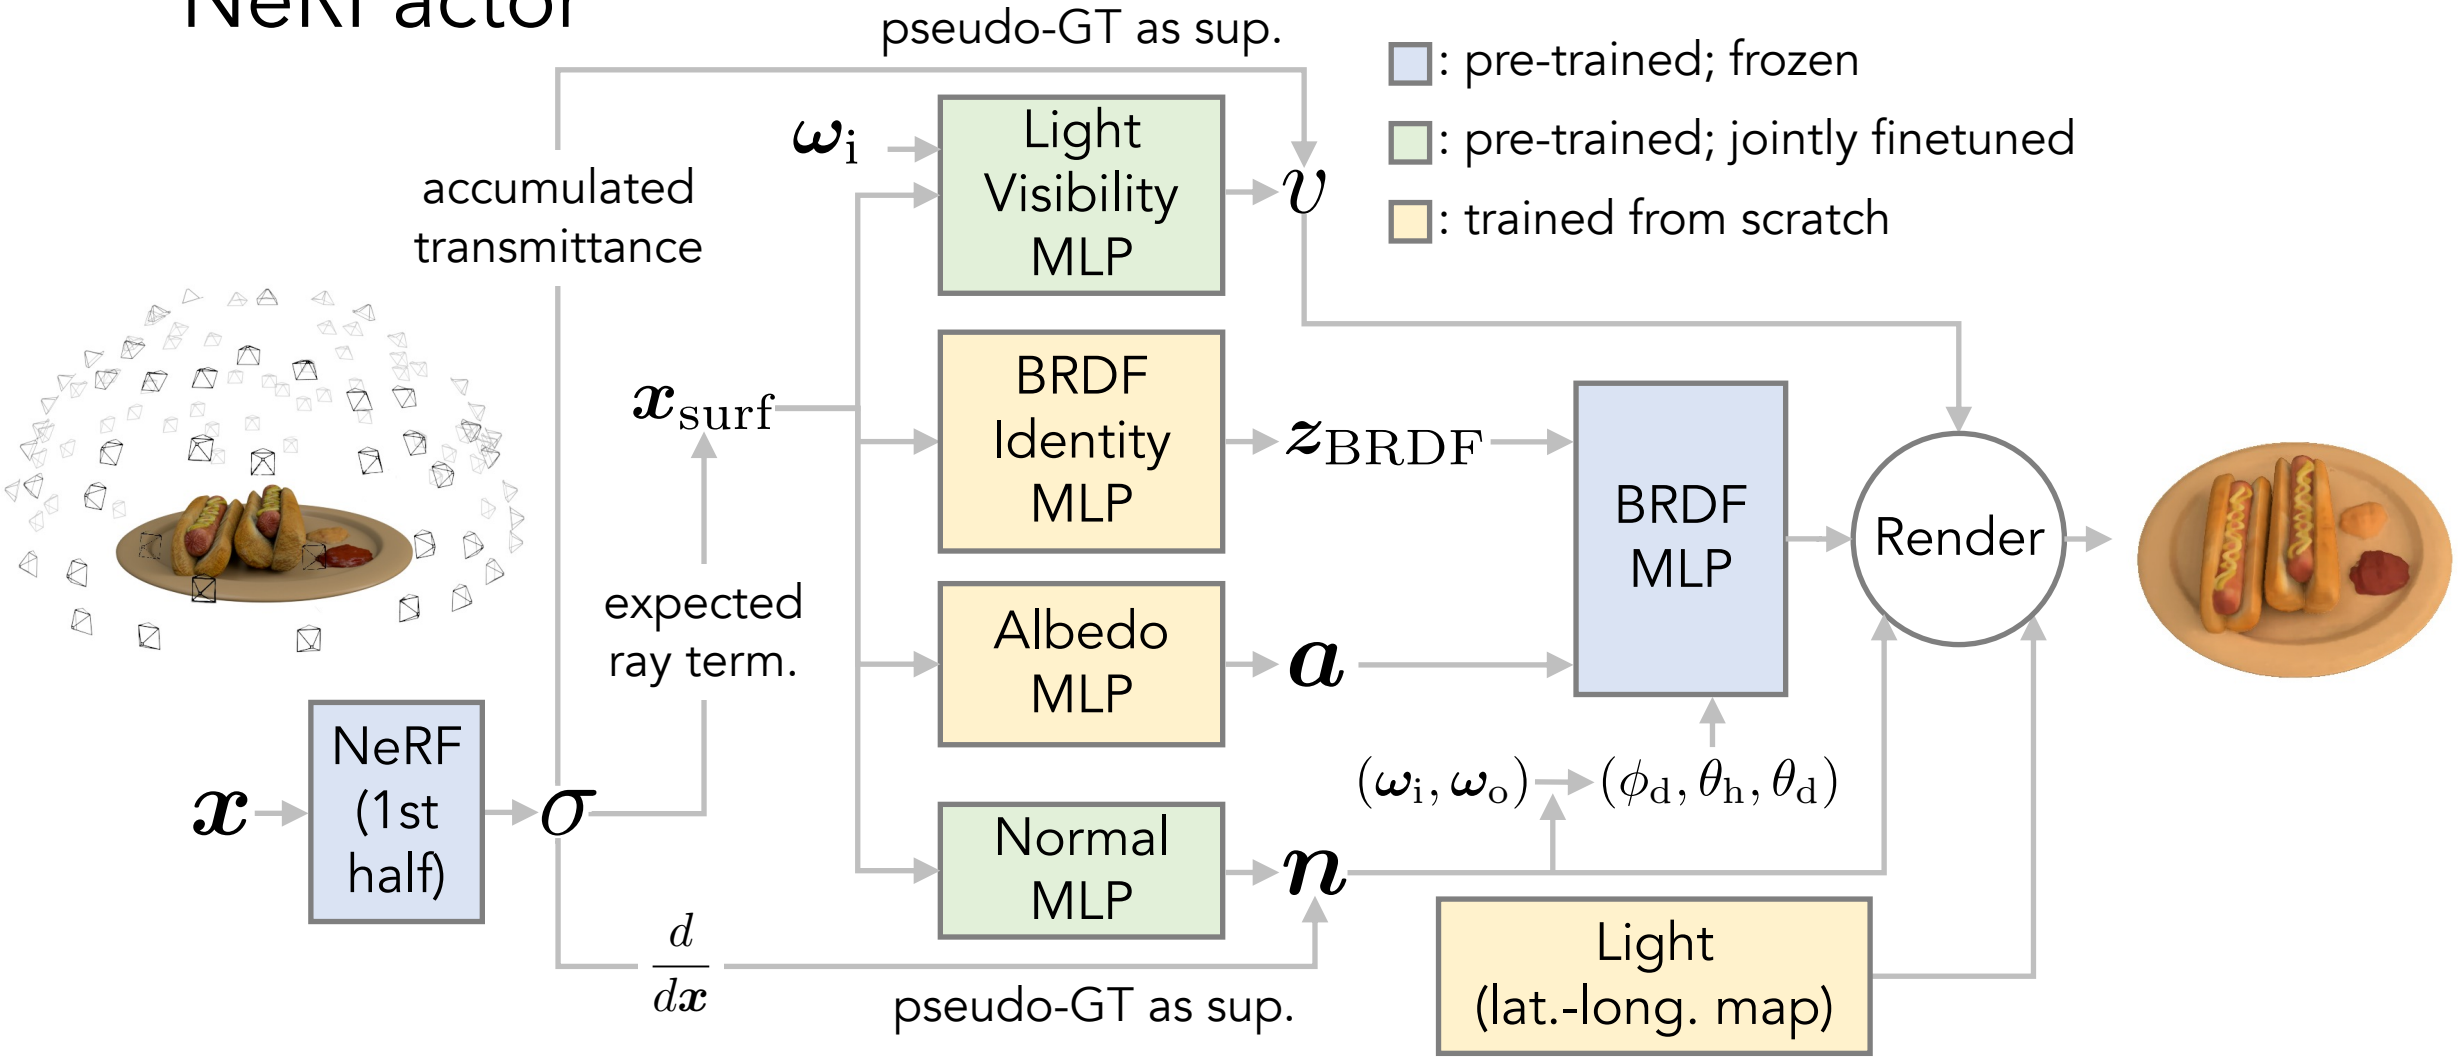

Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, Jonathan T. Barron
NeRFactor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination
arXiv

NeRFactor

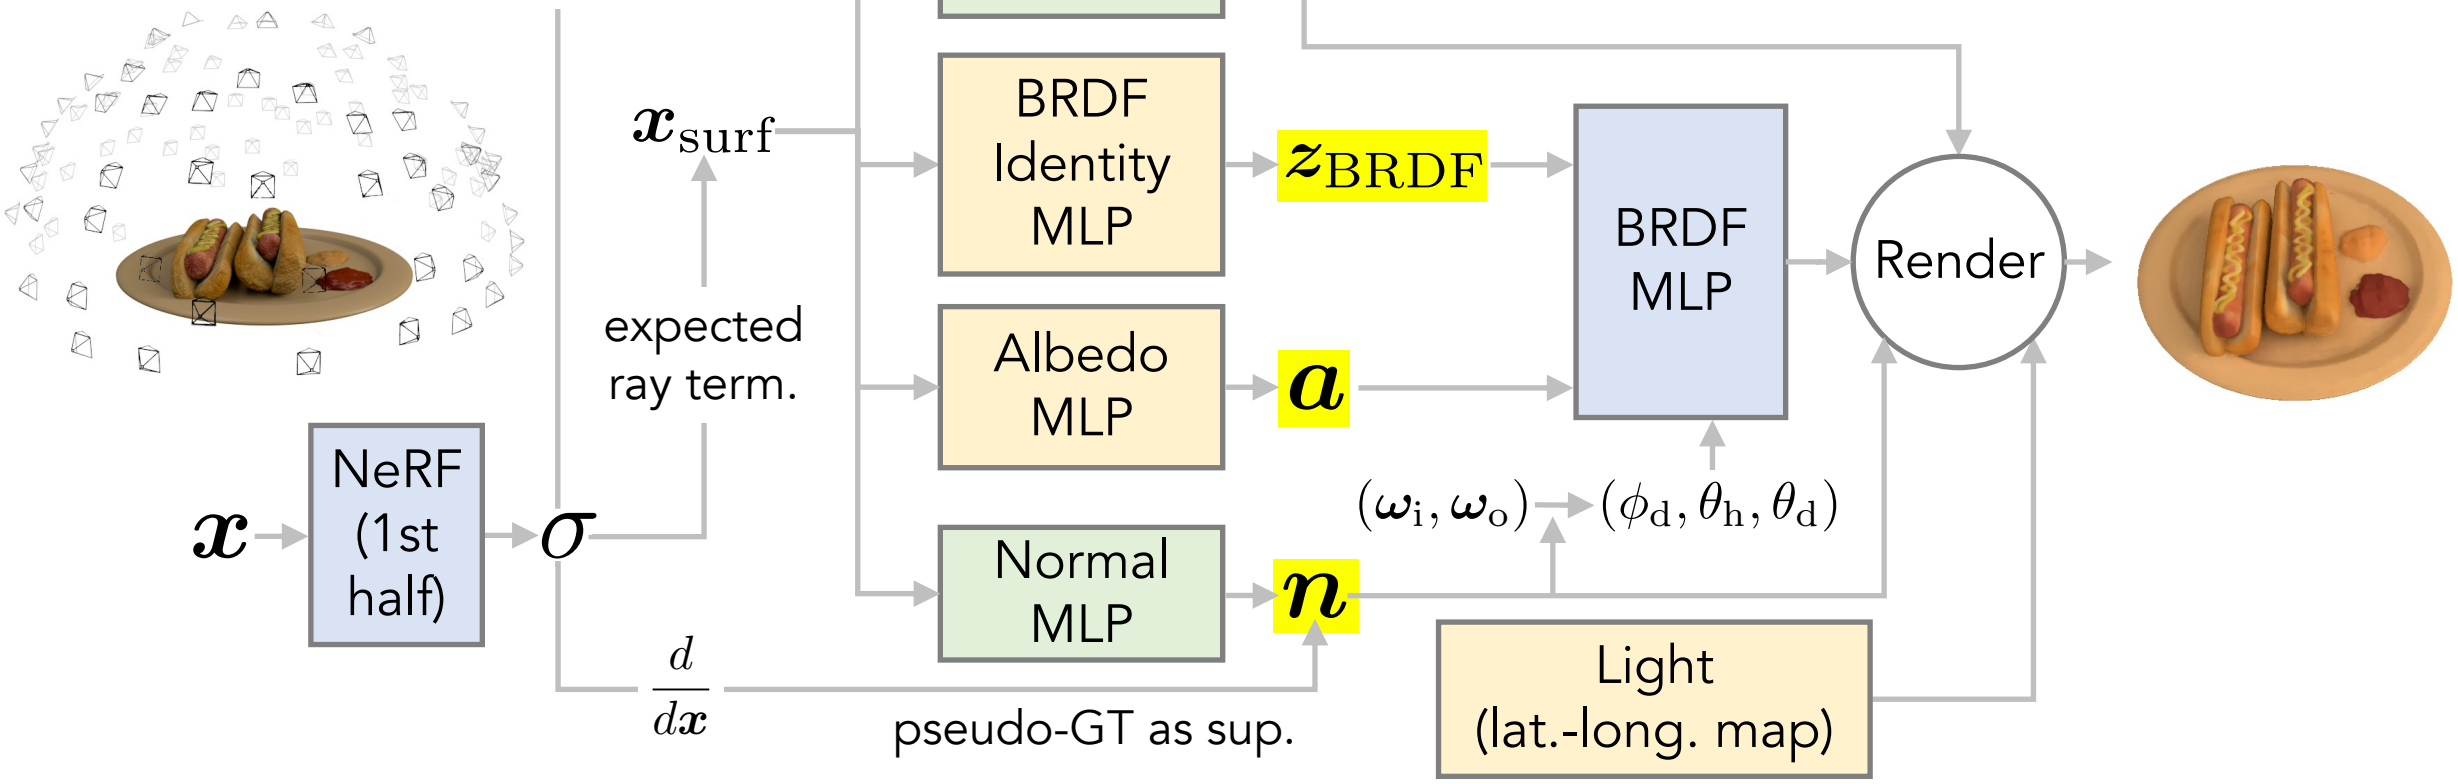

L1 smoothness regularization

Results

on Appearance Factorization, Free-Viewpoint Relighting, & Material Editing

Normals

Visibility

Nearest Input

BRDF

Albedo

Rendering

Normals

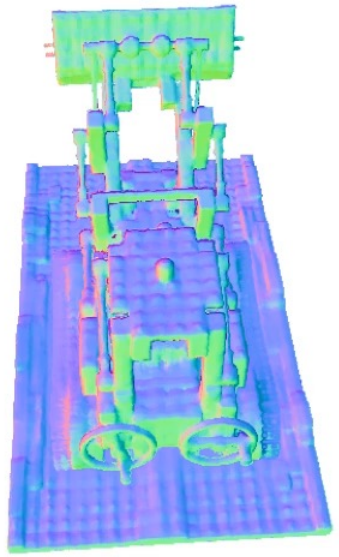

Visibility

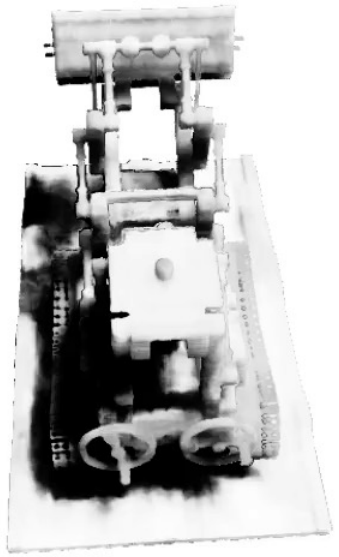

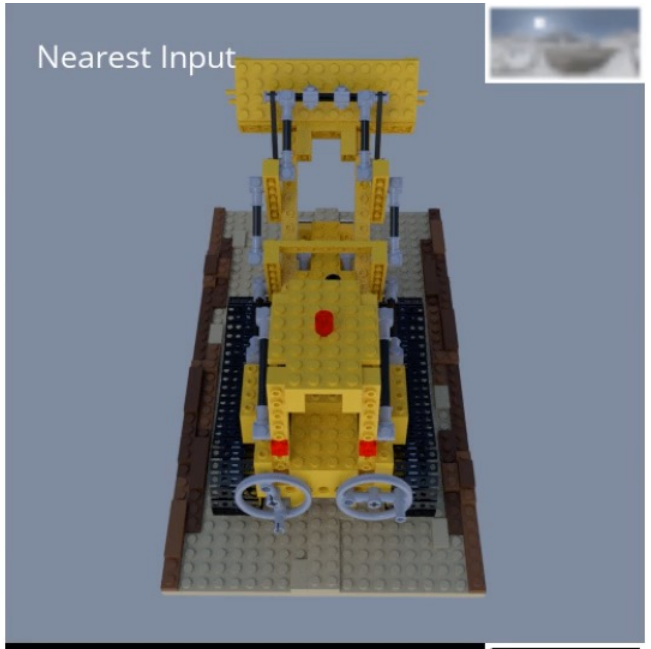

BRDF

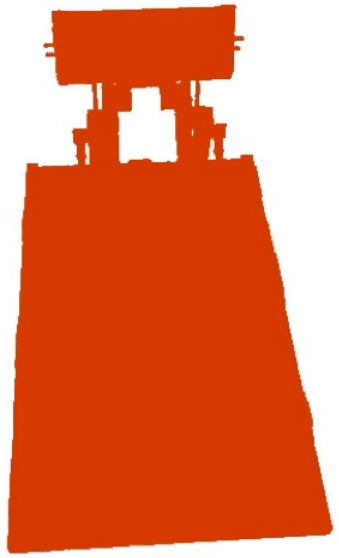

Albedo

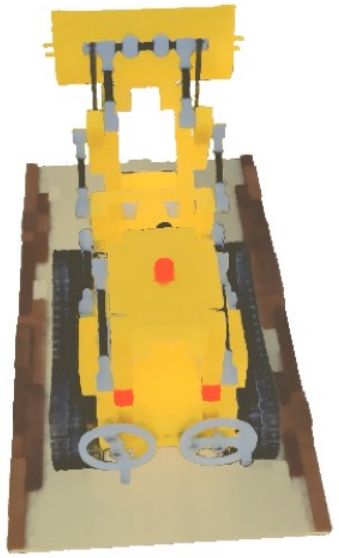

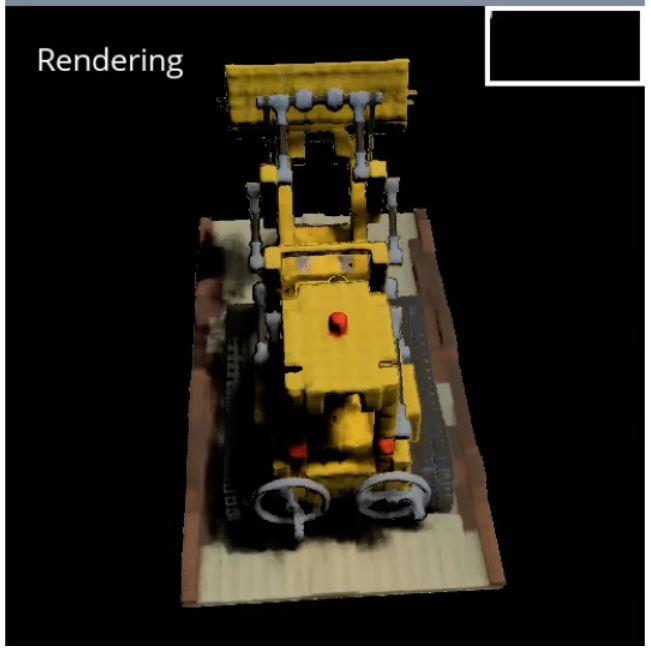

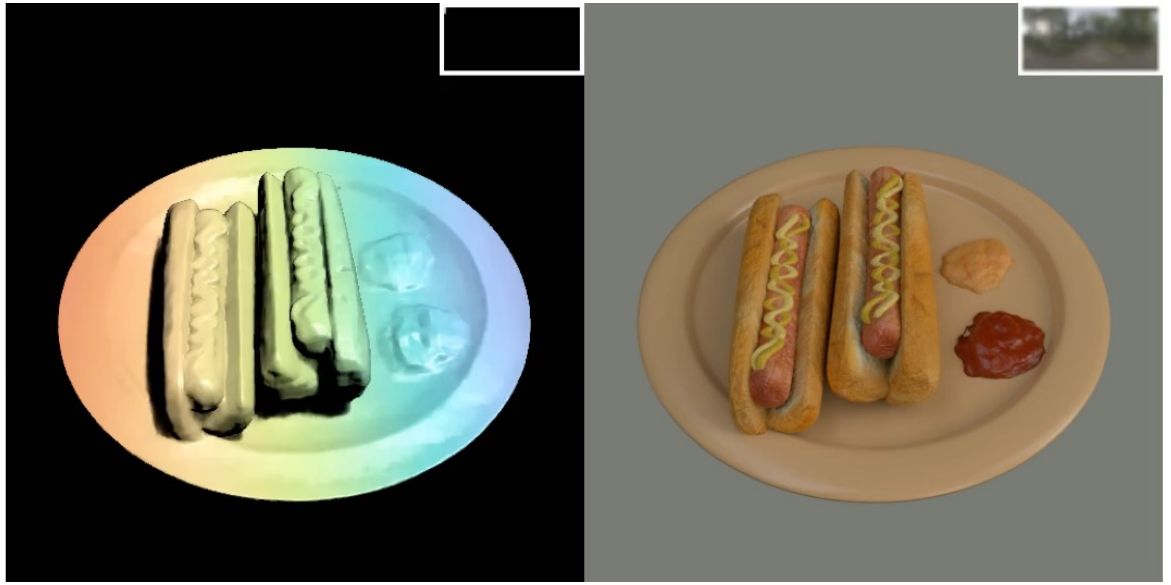

Edited Albedo

Relighting

Nearest Input

Ablation Studies

NeRF’s shape is too noisy for relighting.

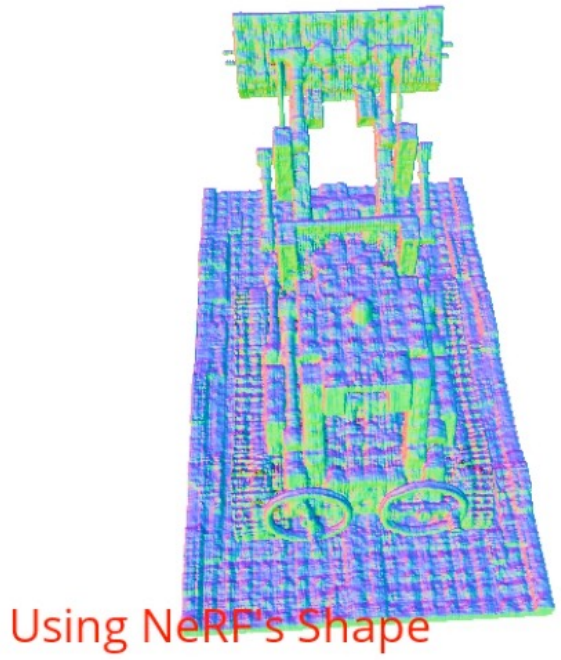

Surface Normals

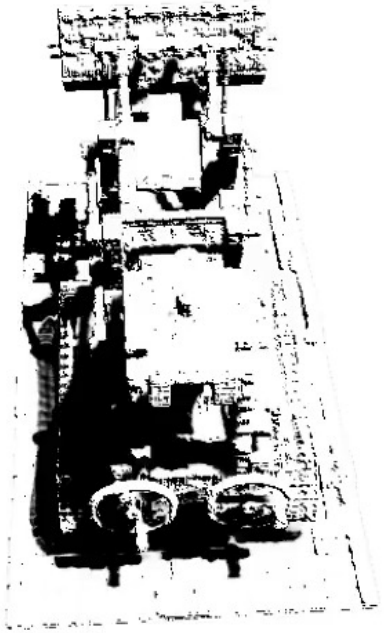

Light Visibility

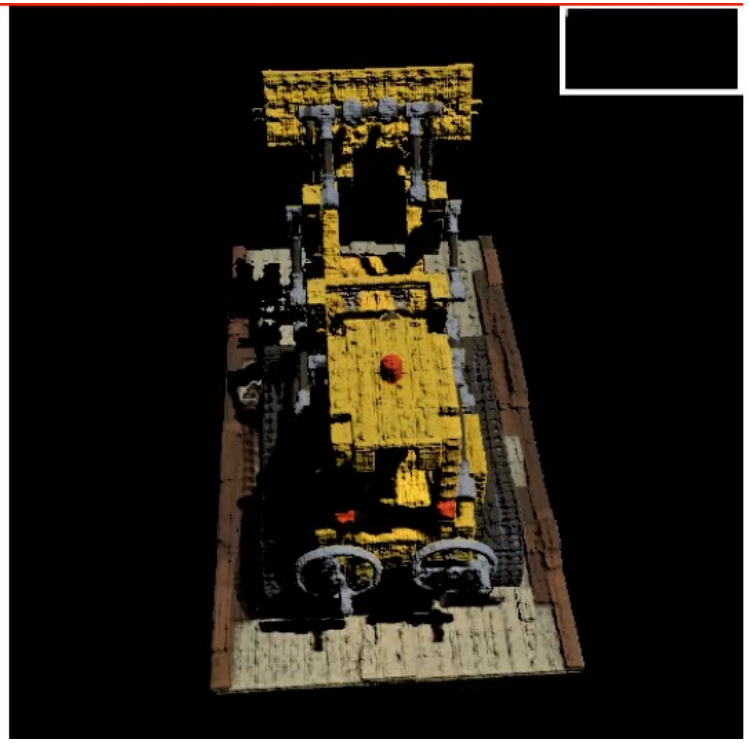

Relighting

Smoothness constraints are crucial.

No Smooth.

Surface Normals

BRDFs

Relighting

Geometry pretraining improves shape and reflectance.

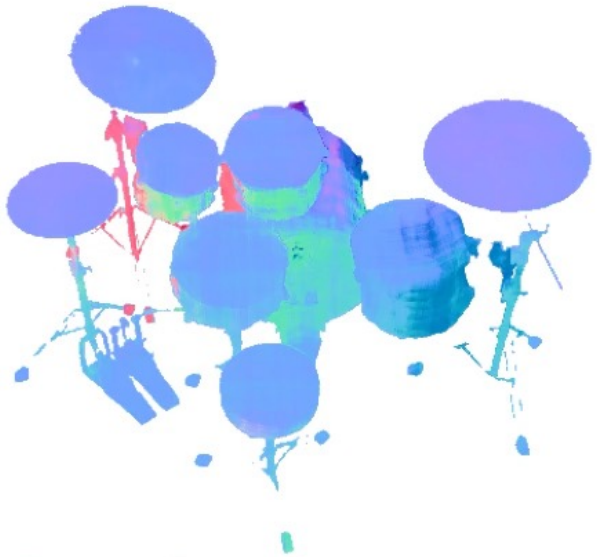

Surface Normals | Albedo | Relighting

Learned BRDFs are more optimization-friendly.

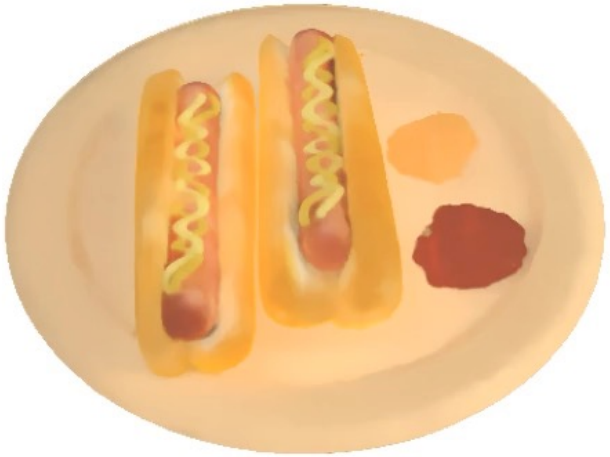

Microfacet

Albedo

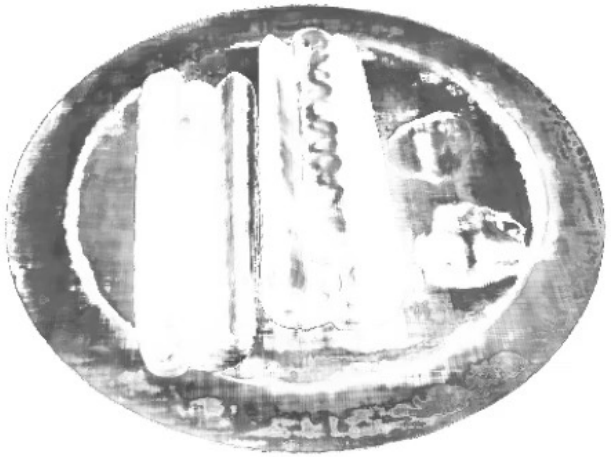

BRDFs

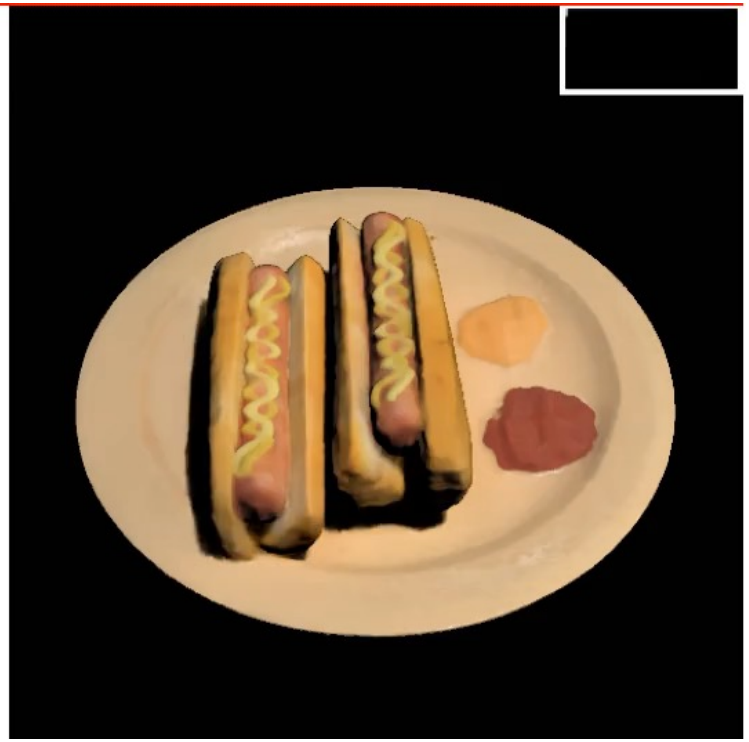

Relighting

Baseline Comparisons

SIRFS does not exploit multiple views.

SIRFS

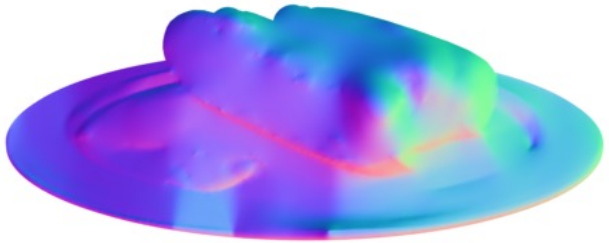

Ground Truth

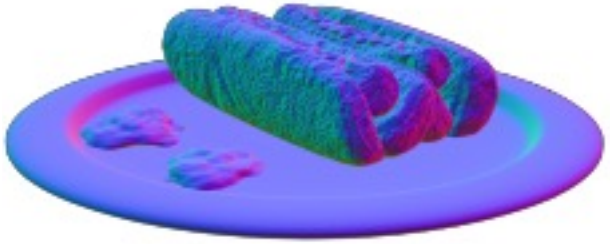

Jonathan T. Barron, Jitendra Malik
Shape, Illumination, and Reflectance From Shading
TPAMI 2015

SIRFS does not exploit multiple views.

NeRFactor (ours)

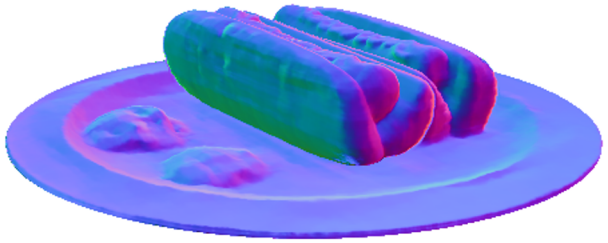

Ground Truth

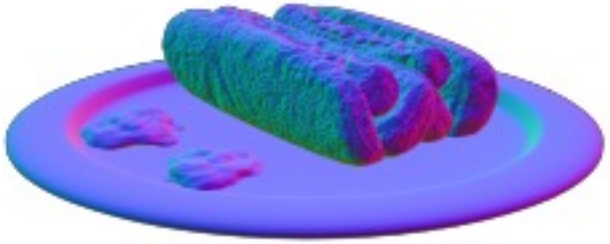

Jonathan T. Barron, Jitendra Malik
Shape, Illumination, and Reflectance From Shading
TPAMI 2015

Oxholm & Nishino do not support SVBRDFs or shadows.

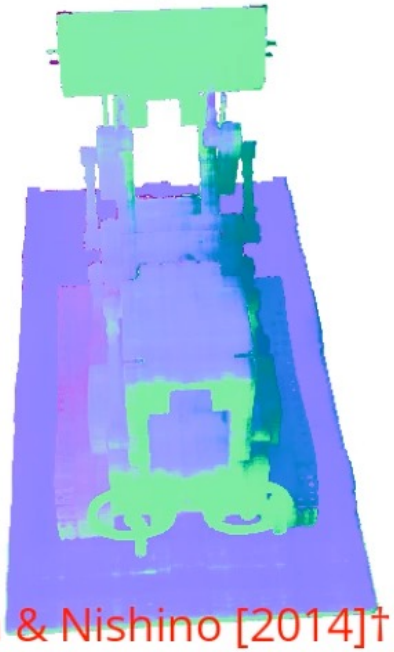

Oxholm & Nishino [2014]†

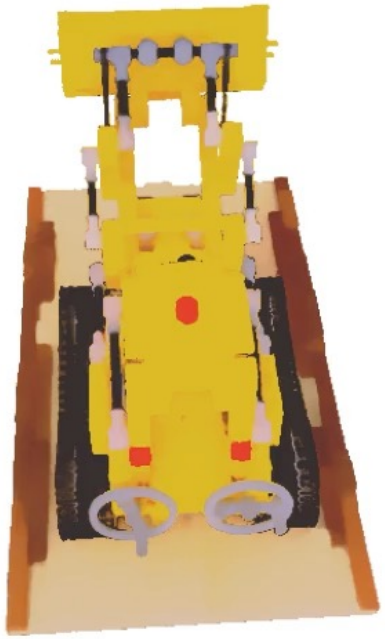

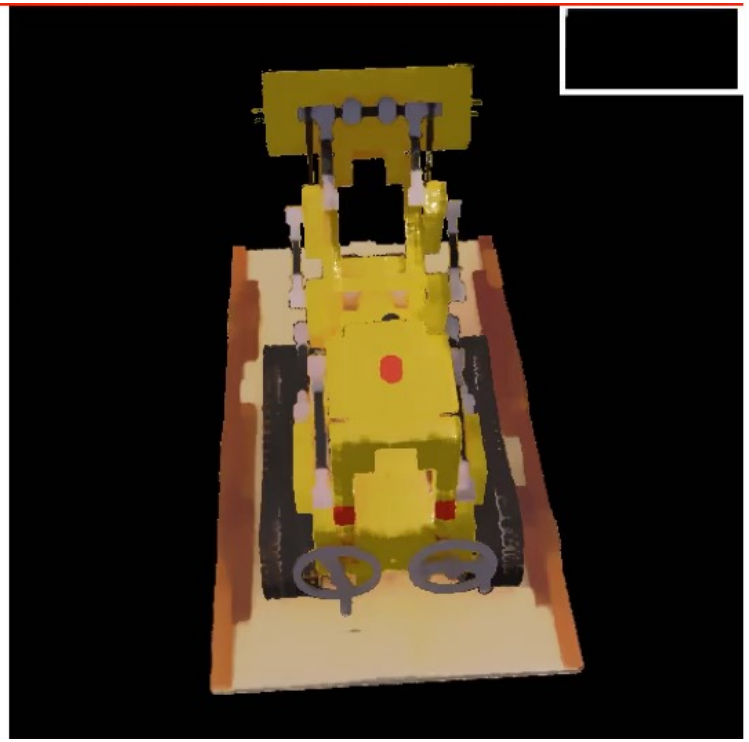

†Enhanced and Requiring Ground-Truth Lighting

Surface Normals

Albedo

Rendering

Geoffrey Oxholm, Ko Nishino
Multiview Shape and Reflectance From Natural Illumination
CVPR 2014

Philip et al. do not support arbitrary relighting.

Philip et al. [2019]

Ground Truth

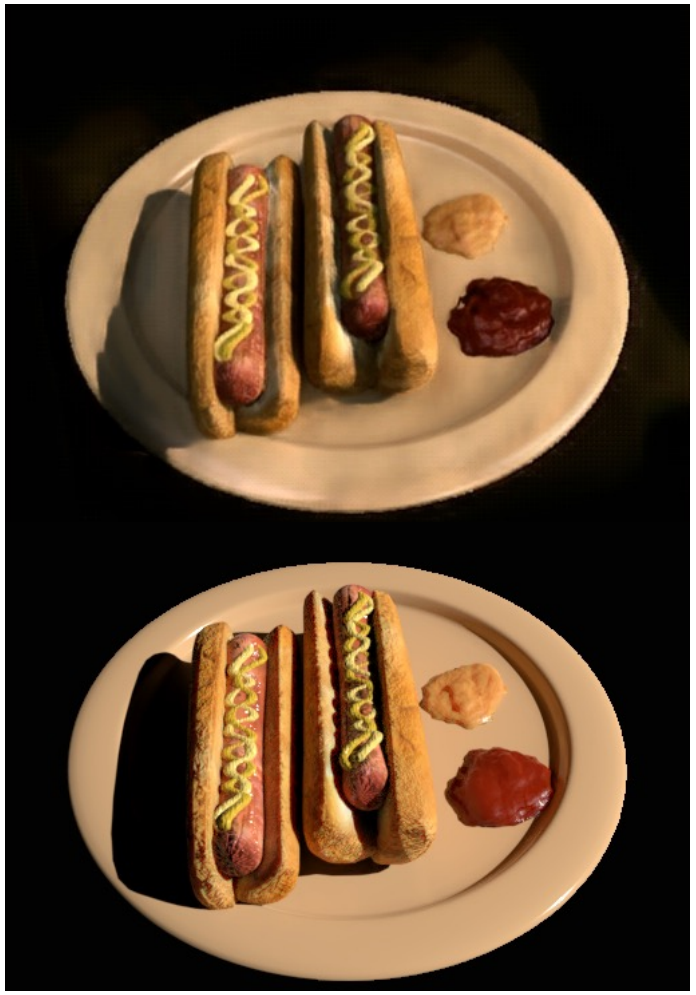

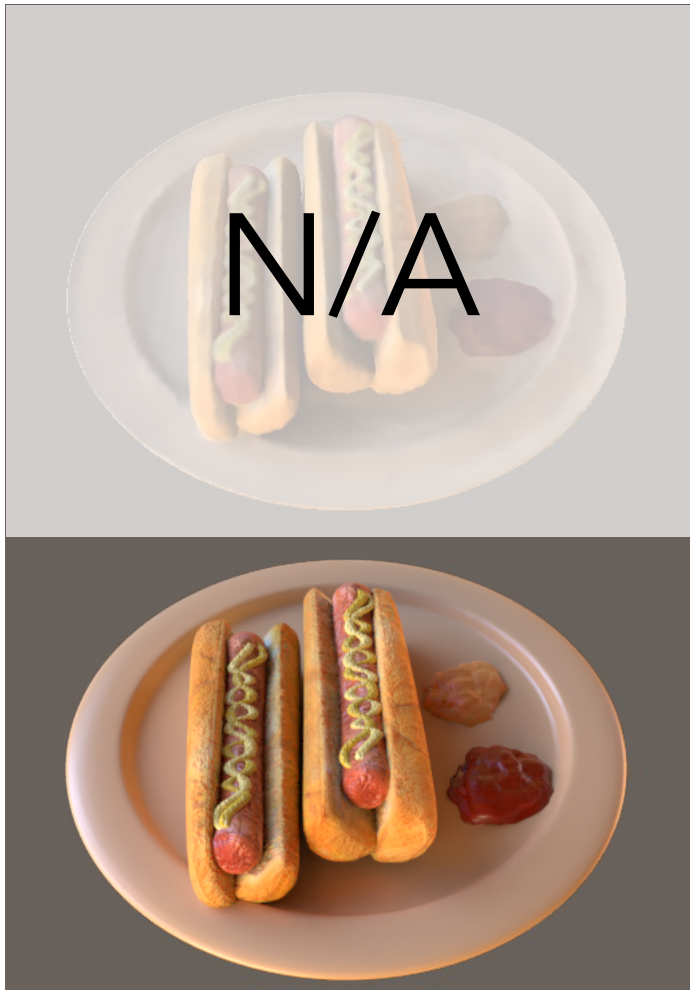

Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A. Efros, George Drettakis
Multi-View Relighting Using a Geometry-Aware Network
TOG 2019

Philip et al. do not support arbitrary relighting.

NeRFactor (ours)

Ground Truth

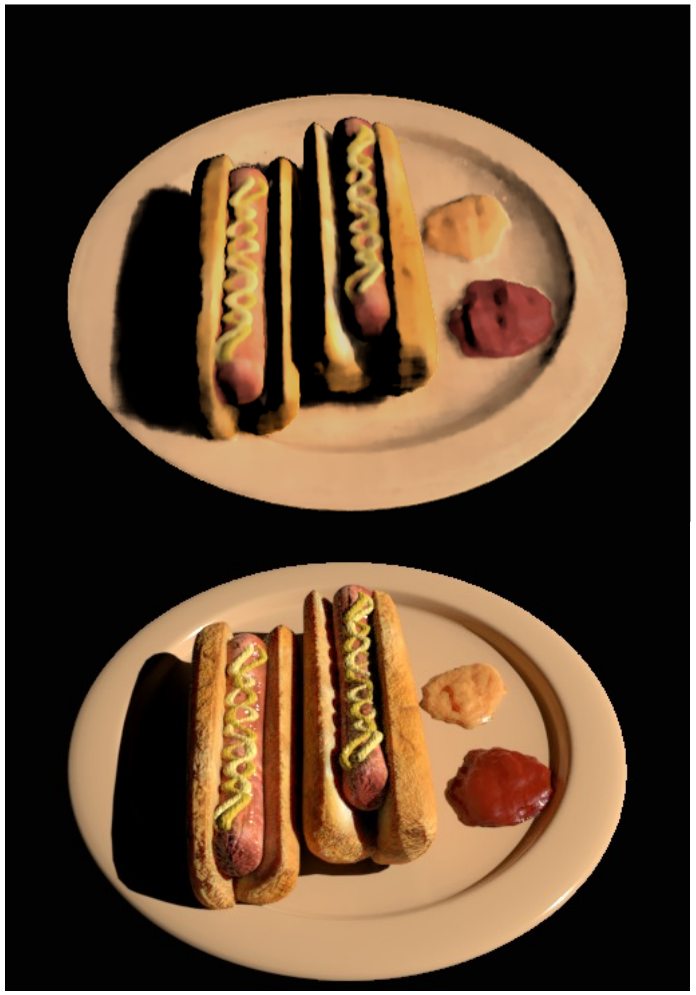

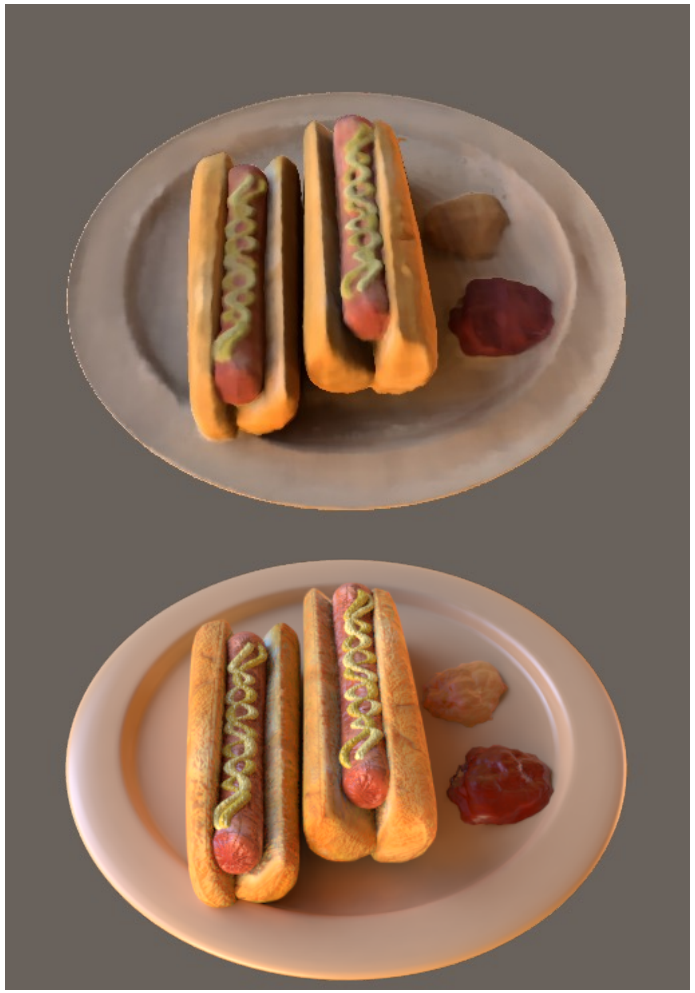

Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A. Efros, George Drettakis
Multi-View Relighting Using a Geometry-Aware Network
TOG 2019

Neural Factorization of Shape and Reflectance (NeRFactor)

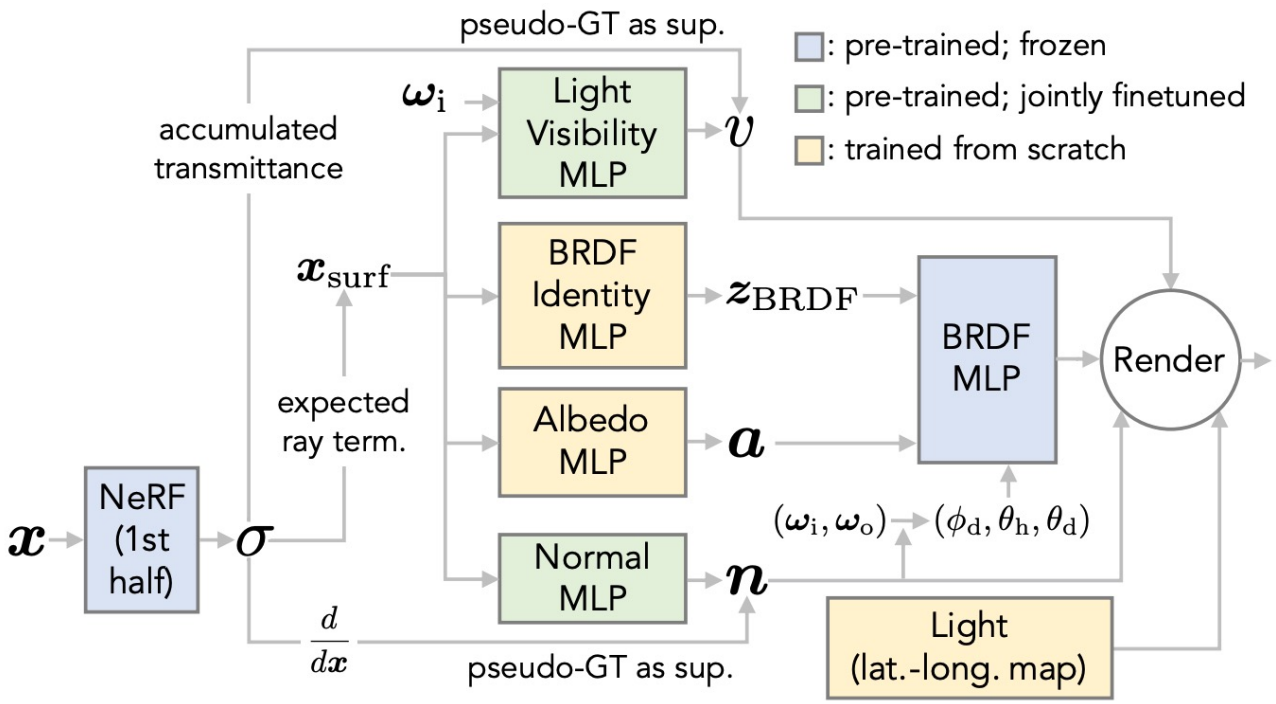

Neural Factorization of Shape and Reflectance (NeRFactor)

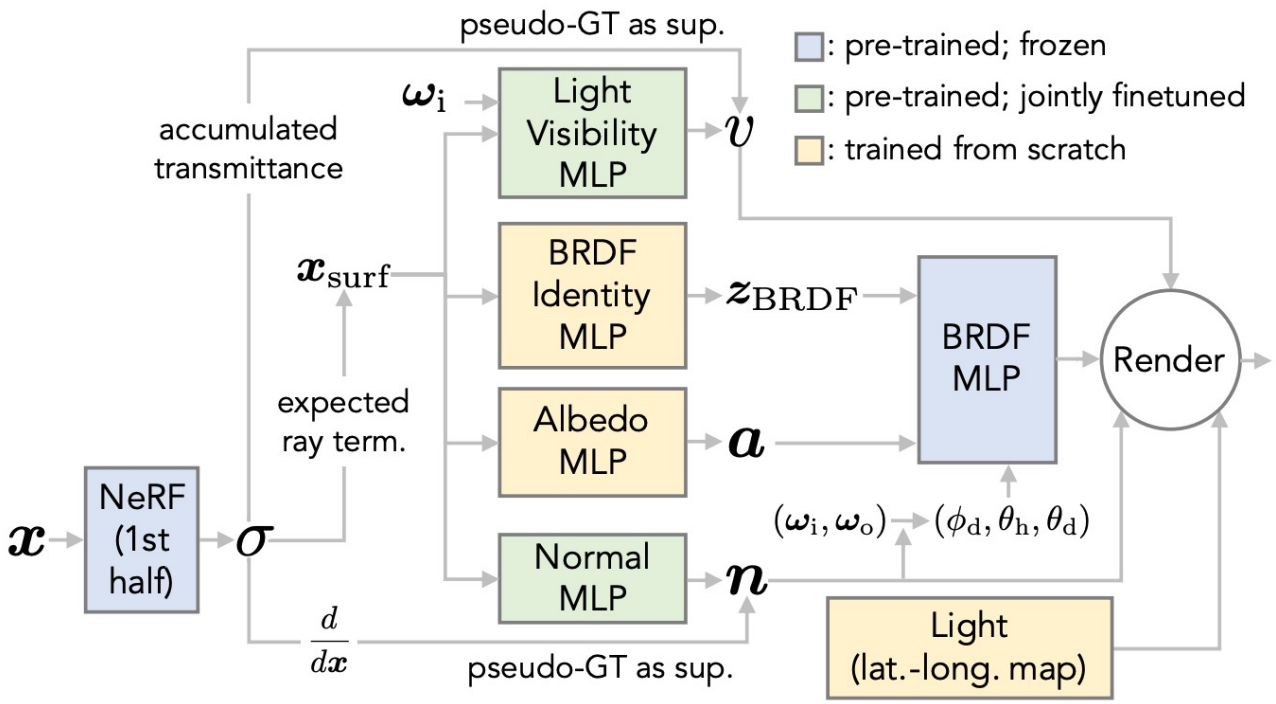

Contributions:

1. A method for factorizing appearance into shape and reflectance, under *one unknown* lighting condition.

Neural Factorization of Shape and Reflectance (NeRFactor)

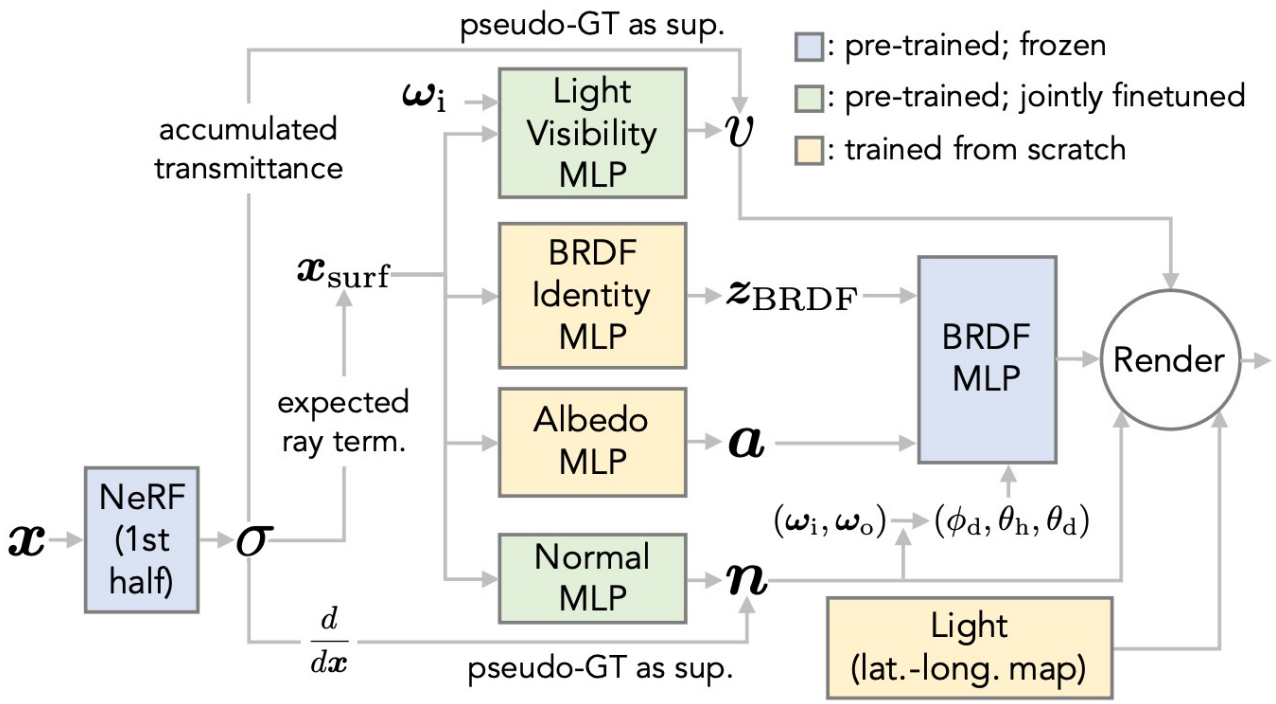

Contributions:

1. A method for factorizing appearance into shape and reflectance, under *one unknown* lighting condition.
2. A strategy to distill NeRF-estimated geometry and then refine it.

Neural Factorization of Shape and Reflectance (NeRFactor)

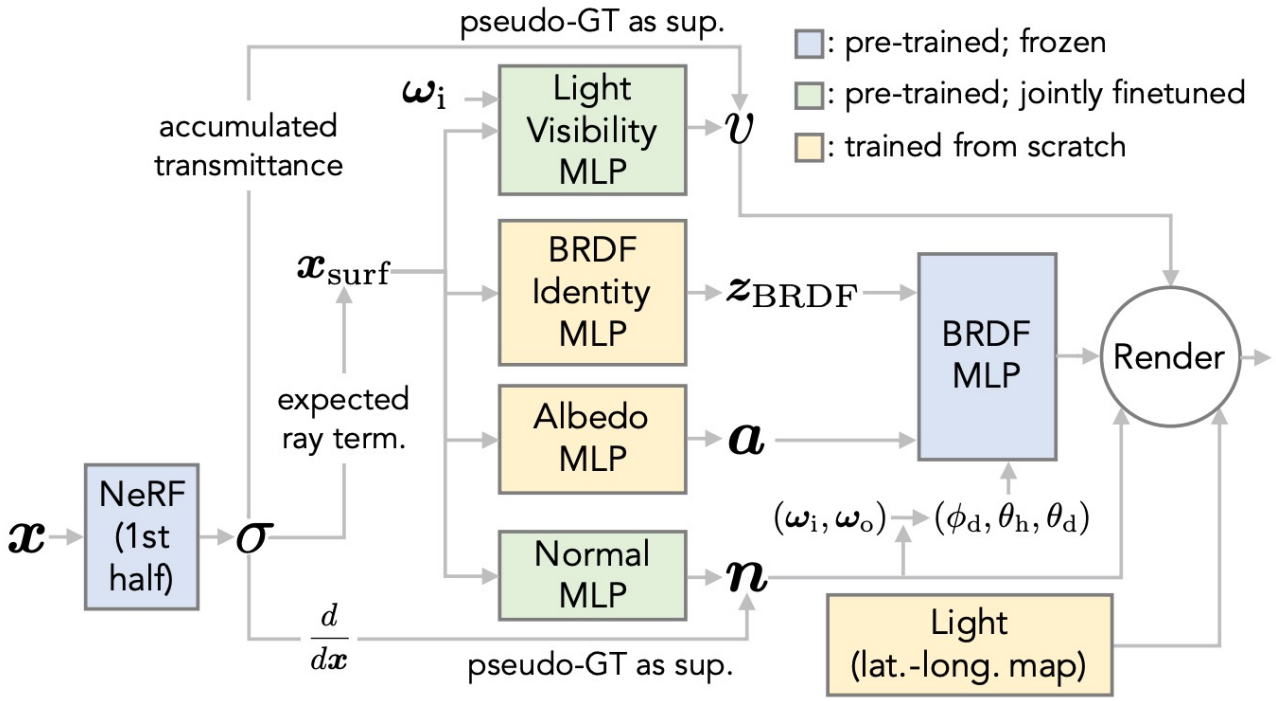

Contributions:

1. A method for factorizing appearance into shape and reflectance, under *one unknown* lighting condition.
2. A strategy to distill NeRF-estimated geometry and then refine it.
3. Data-driven BRDF priors learned from real measured BRDFs.

Neural Factorization of Shape and Reflectance (NeRFactor)

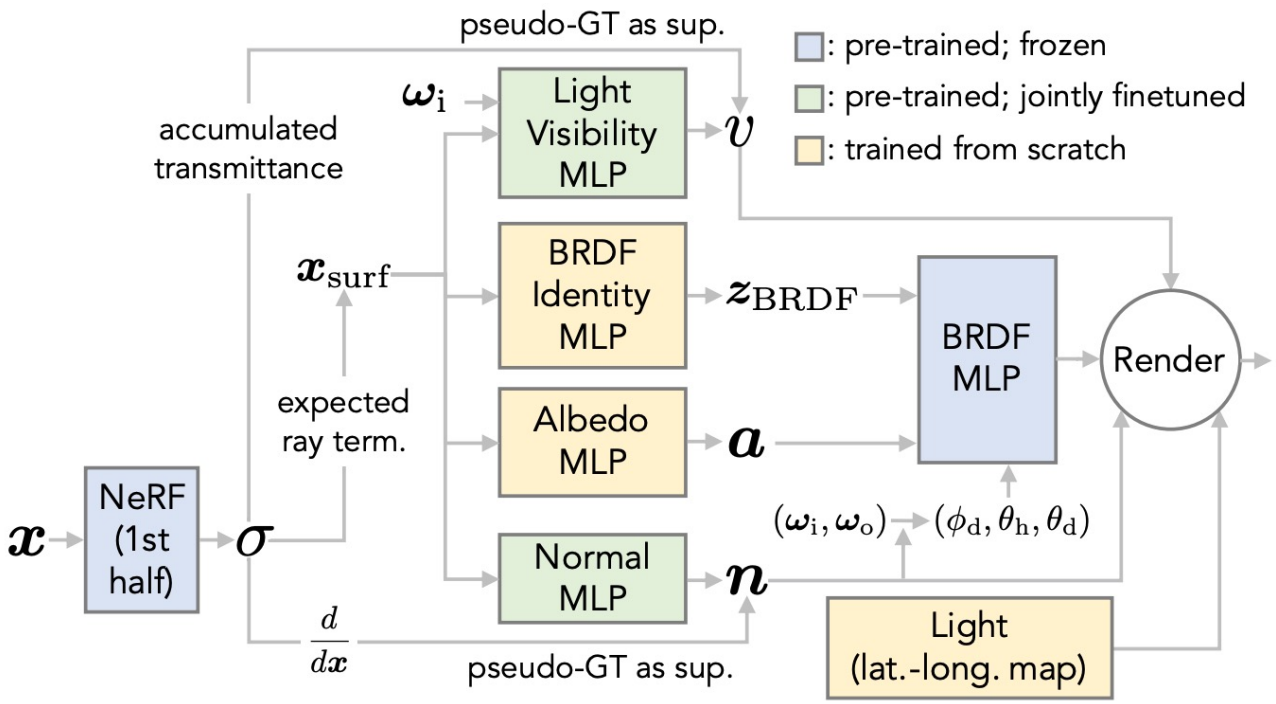

Contributions:

1. A method for factorizing appearance into shape and reflectance, under *one unknown* lighting condition.
2. A strategy to distill NeRF-estimated geometry and then refine it.
3. Data-driven BRDF priors learned from real measured BRDFs.

❌ No indirect illumination

❌ Relies on NeRF geometry

Neural Factorization of Shape and Reflectance (NeRFactor)

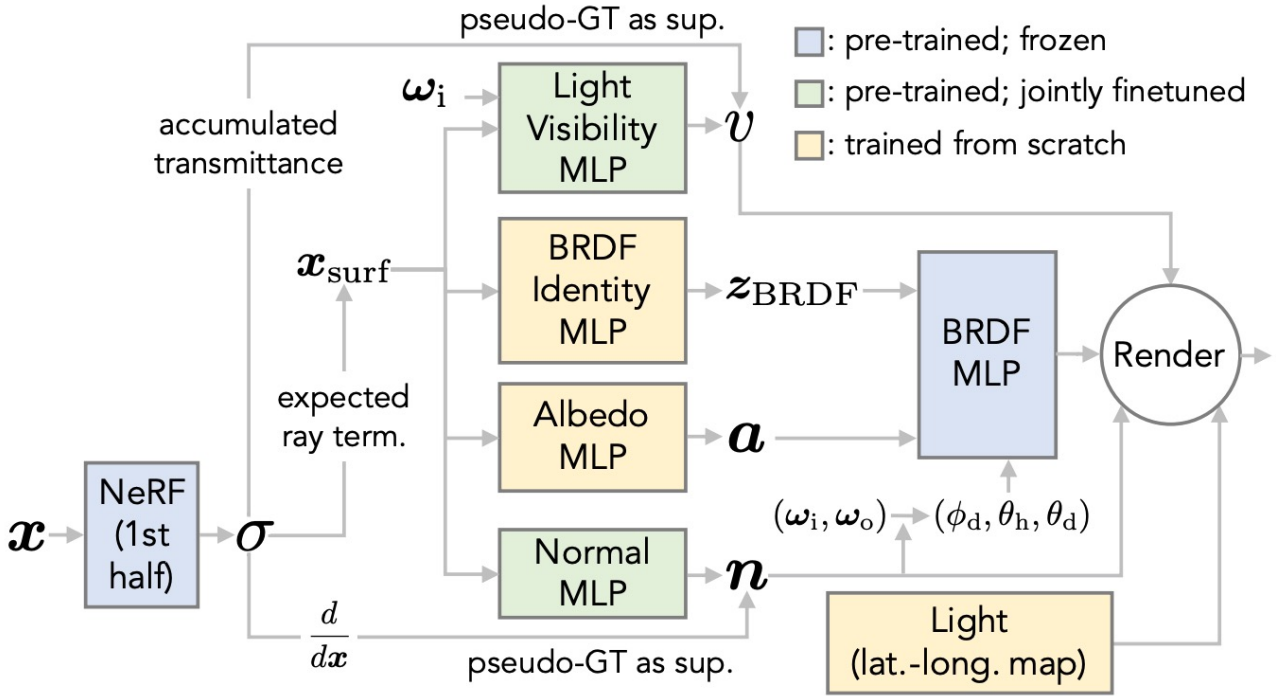

Contributions:

1. A method for factorizing appearance into shape and reflectance, under *one unknown* lighting condition.
2. A strategy to distill NeRF-estimated geometry and then refine it.
3. Data-driven BRDF priors learned from real measured BRDFs.

❌ No indirect illumination

❌ Relies on NeRF geometry

Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, Jonathan T. Barron

NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis

CVPR 2021

Neural Factorization of Shape and Reflectance (NeRFactor)

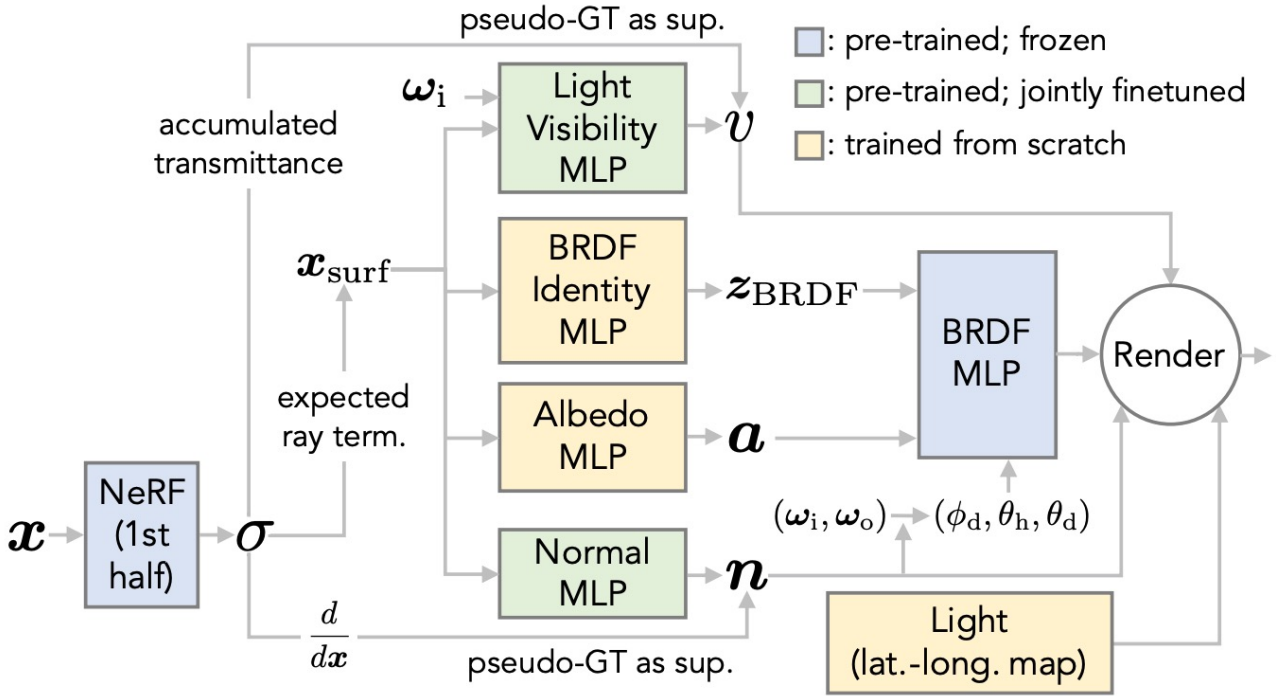

Contributions:

1. A method for factorizing appearance into shape and reflectance, under *one unknown* lighting condition.
2. A strategy to distill NeRF-estimated geometry and then refine it.
3. Data-driven BRDF priors learned from real measured BRDFs.

❌ No indirect illumination

❌ Relies on NeRF geometry

Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, Jonathan T. Barron

NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis

CVPR 2021

Key differences:

1. Models indirect illumination,
2. Optimizes shape from scratch, and
3. Learns from multiple, known lighting conditions.

Neural Visibility Fields (NeRV)

(a) Our Rendered Image
(Novel View and Lighting)

Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, Jonathan T. Barron
NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis
CVPR 2021

Neural Visibility Fields (NeRV)

(a) Our Rendered Image
(Novel View and Lighting)

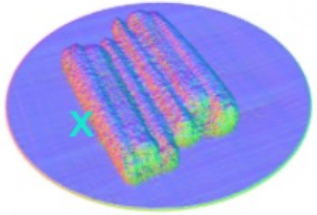

(f) Normals

Neural Visibility Fields (NeRV)

$$= \int_S ($$

(b) Light Visibility

(a) Our Rendered Image
(Novel View and Lighting)

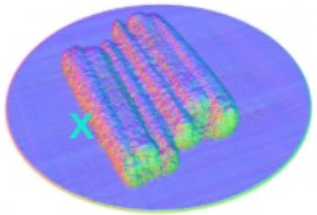

(f) Normals

Neural Visibility Fields (NeRV)

(a) Our Rendered Image
(Novel View and Lighting)

$$= \int_{\mathcal{S}} ($$

(b) Light Visibility

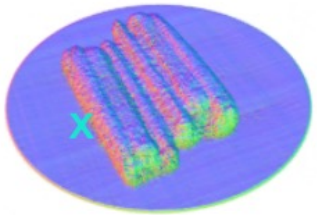

(f) Normals

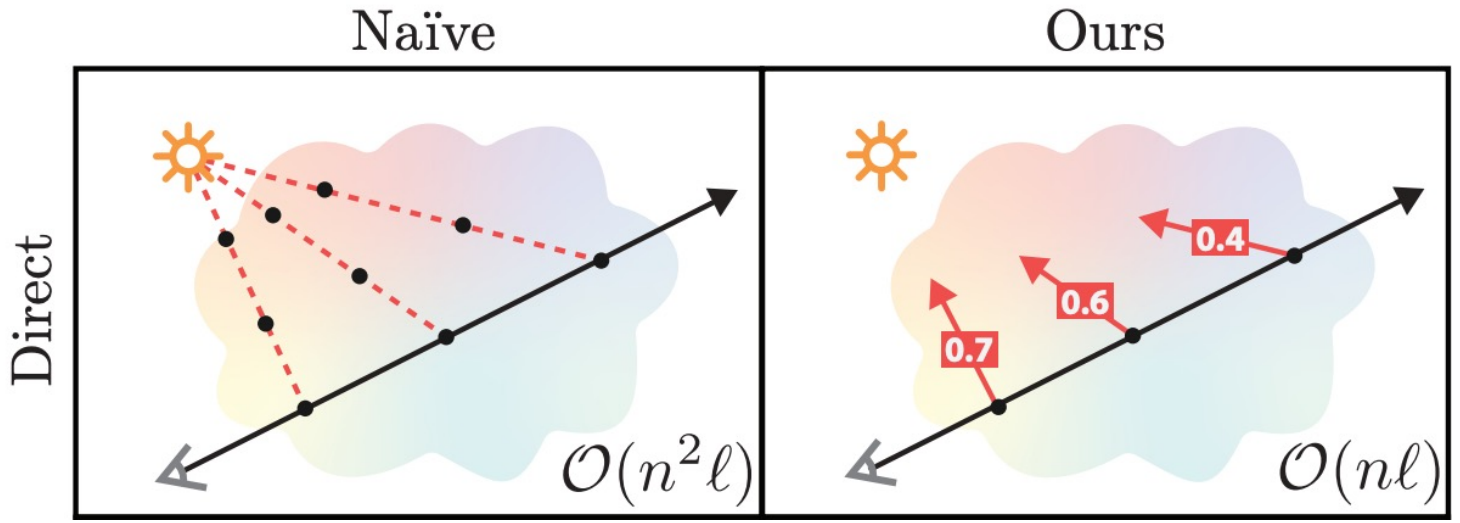

Neural Visibility Fields (NeRV)

$$= \int_S ($$

(b) Light Visibility

(a) Our Rendered Image
(Novel View and Lighting)

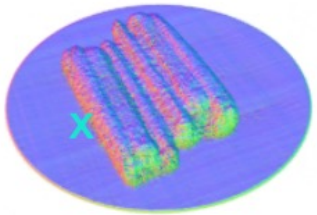

(f) Normals

Neural Visibility Fields (NeRV)

(a) Our Rendered Image
(Novel View and Lighting)

$$= \int_S ($$

(b) Light Visibility

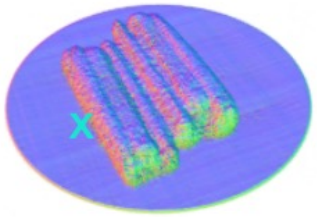

(f) Normals

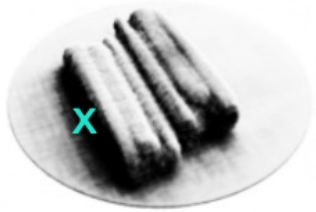

(i) Shadow Map

Neural Visibility Fields (NeRV)

$$= \int_S \Big($$

(b) Light Visibility

(a) Our Rendered Image
(Novel View and Lighting)

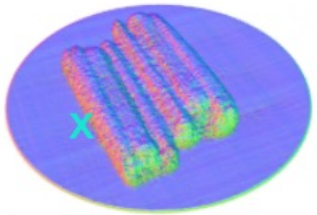

(f) Normals

(g) Albedo

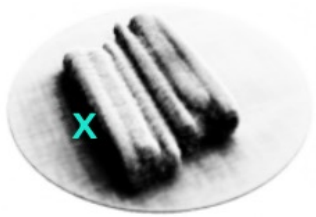

(i) Shadow Map

Neural Visibility Fields (NeRV)

(a) Our Rendered Image
(Novel View and Lighting)

$$= \int_S ($$

(b) Light Visibility

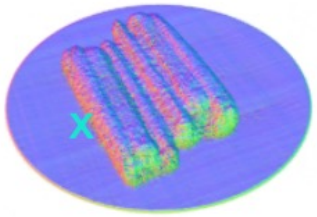

(f) Normals

(g) Albedo

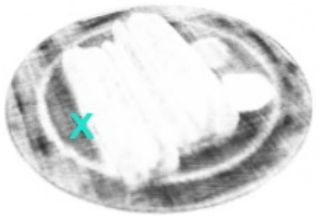

(h) Roughness

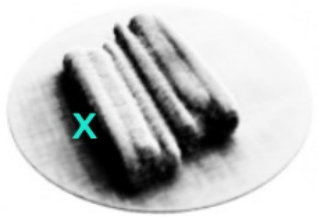

(i) Shadow Map

Neural Visibility Fields (NeRV)

(a) Our Rendered Image
(Novel View and Lighting)

$$= \int_{\mathcal{S}} ($$

(b) Light Visibility

$$) \times$$

(e) BRDF

$$d\omega_i$$

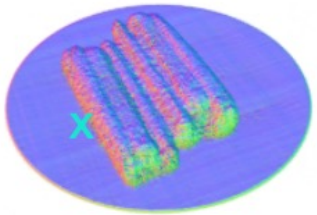

(f) Normals

(g) Albedo

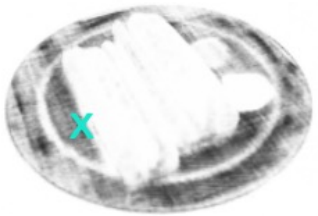

(h) Roughness

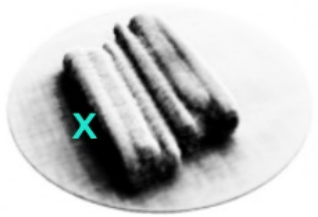

(i) Shadow Map

Neural Visibility Fields (NeRV)

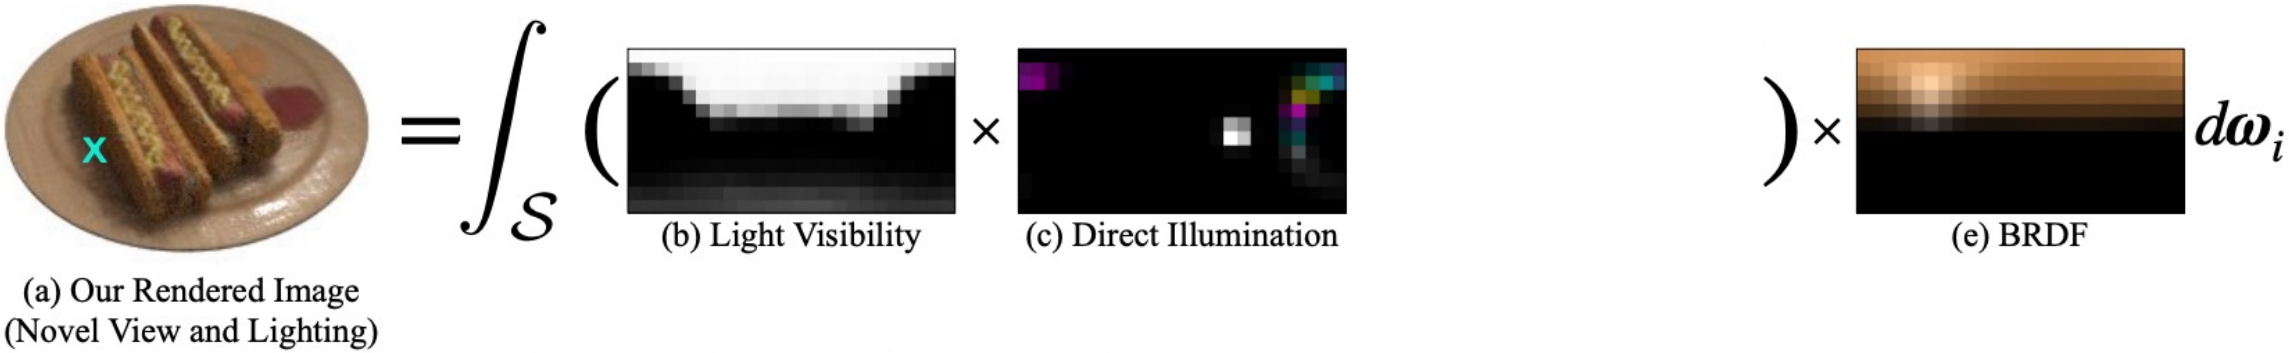

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(e) BRDF

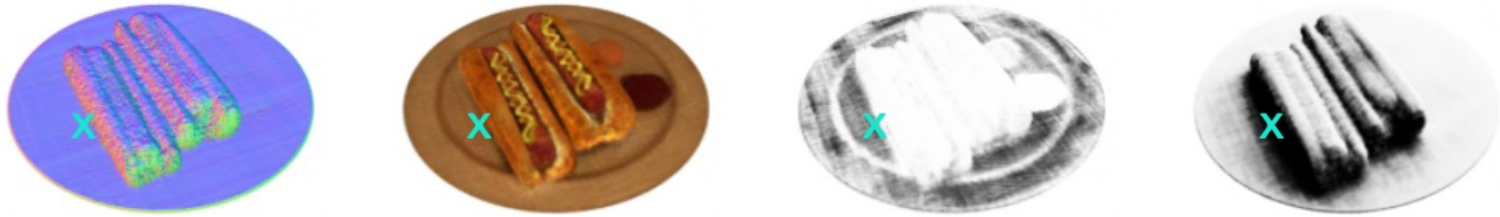

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

Neural Visibility Fields (NeRV)

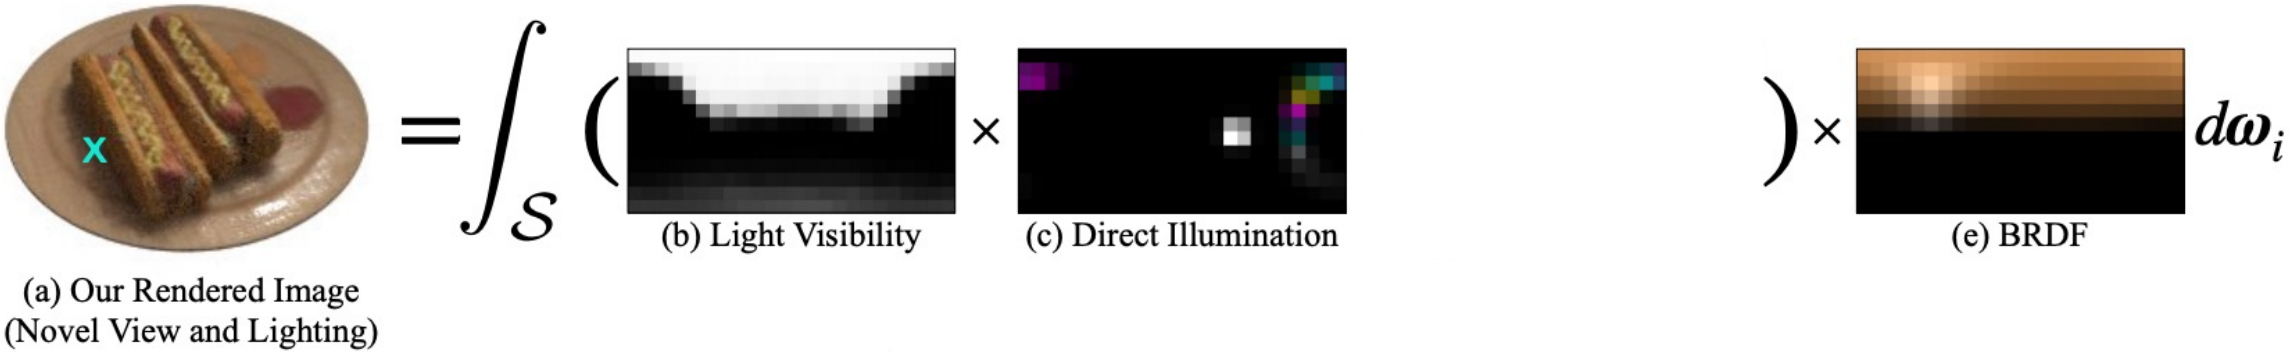

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(e) BRDF

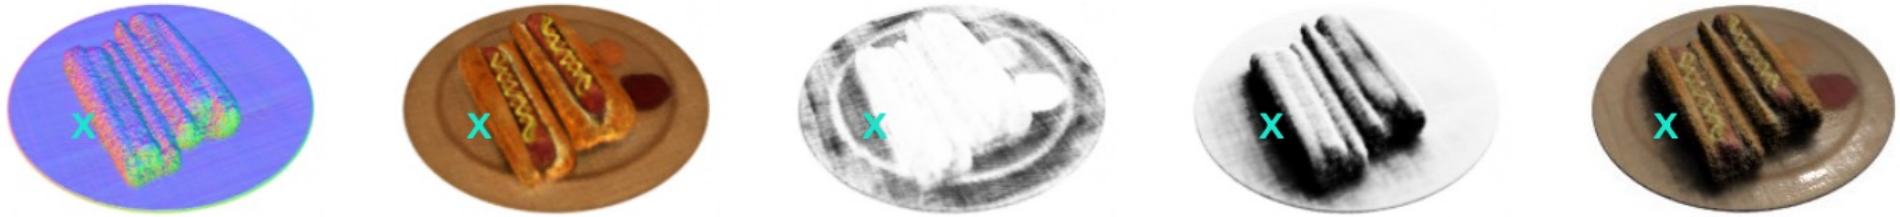

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

(j) Direct

Neural Visibility Fields (NeRV)

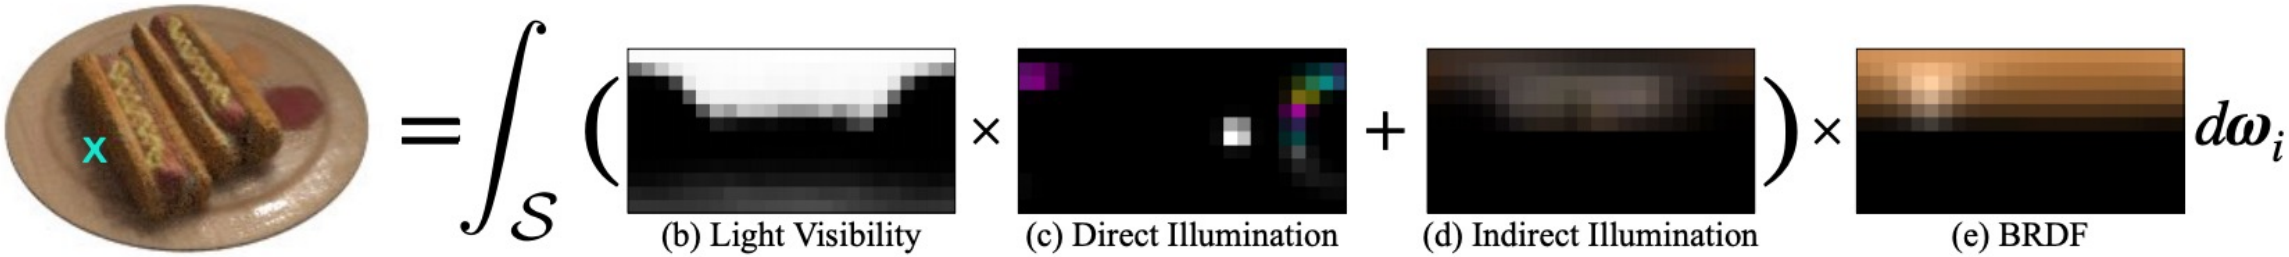

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(d) Indirect Illumination

(e) BRDF

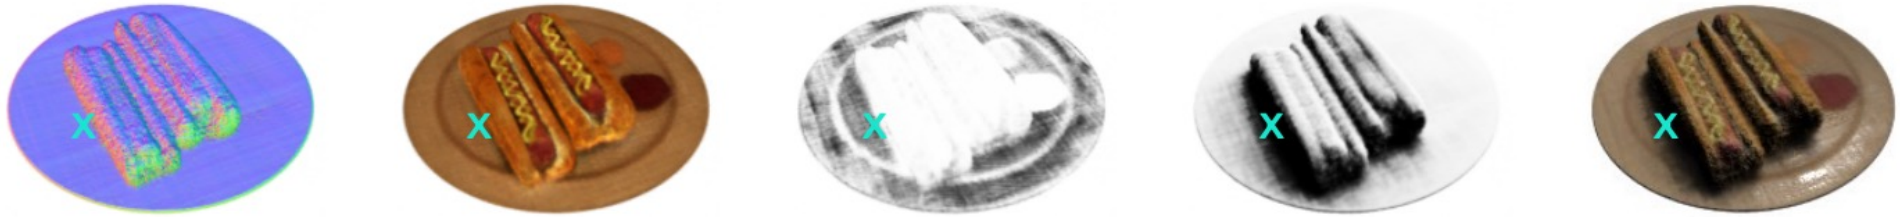

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

(j) Direct

Neural Visibility Fields (NeRV)

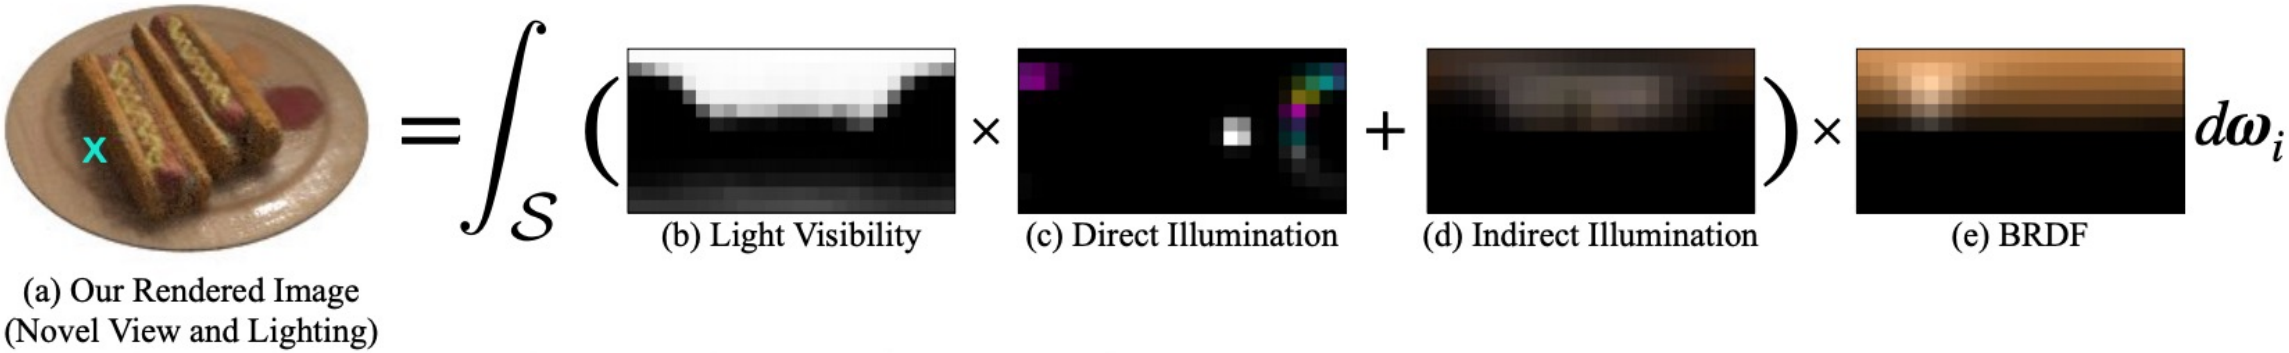

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(d) Indirect Illumination

(e) BRDF

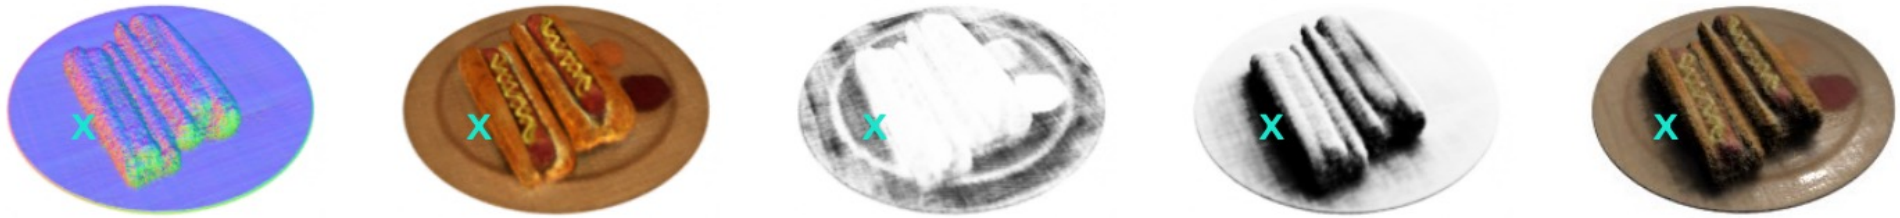

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

(j) Direct

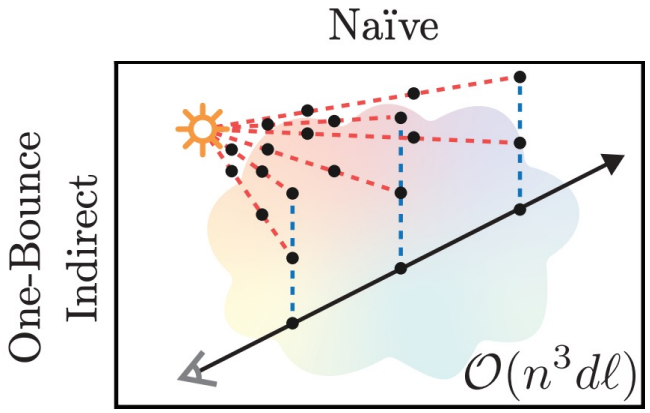

Neural Visibility Fields (NeRV)

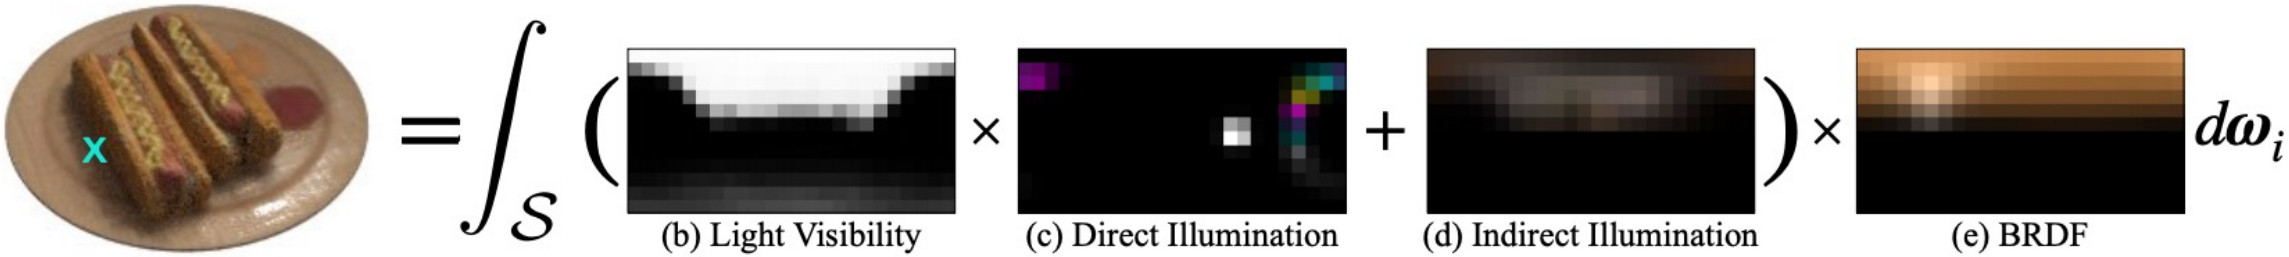

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(d) Indirect Illumination

(e) BRDF

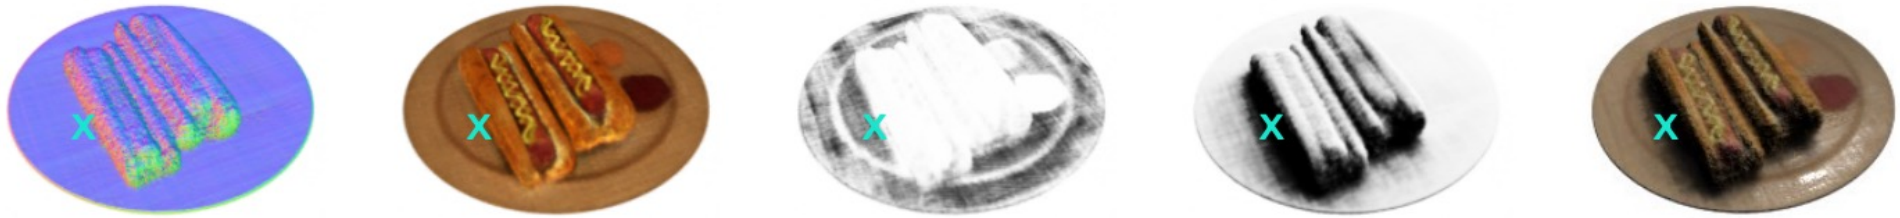

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

(j) Direct

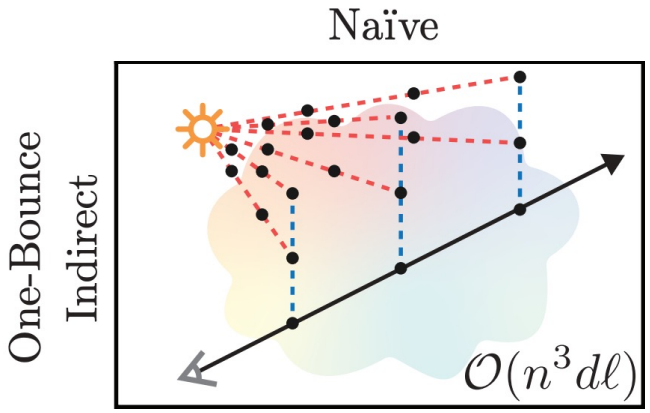

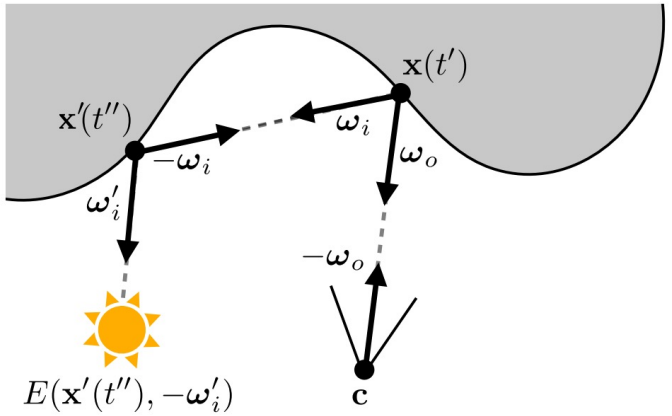

$\mathbf{c}$ = camera center

E = light source

t' = expected termination depth along camera ray

$t'' = \tilde{D}_\phi(\mathbf{x}(t'), \omega_i)$, estimated termination depth along indirect bounce ray

Neural Visibility Fields (NeRV)

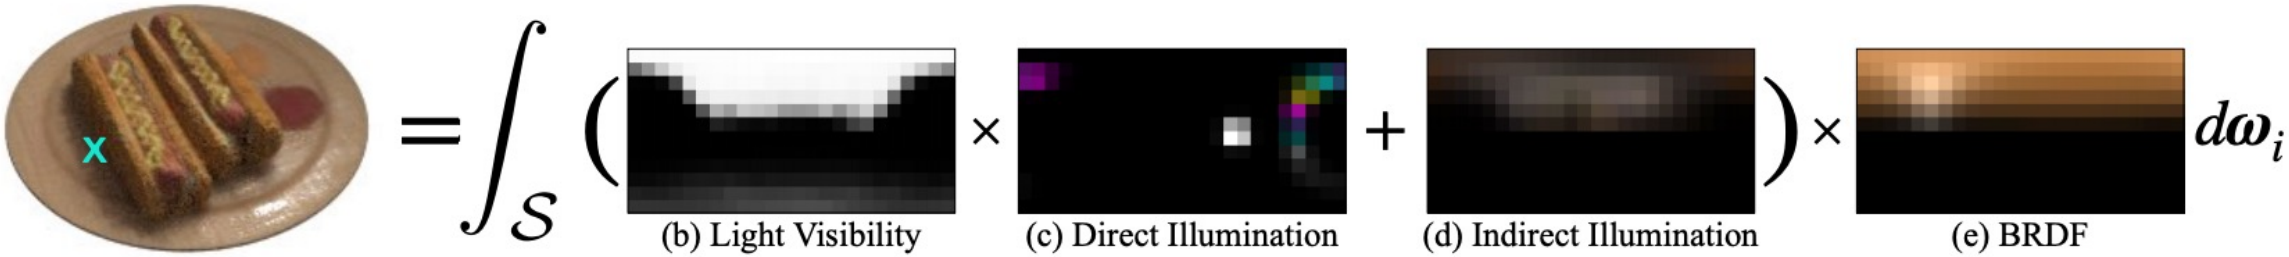

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(d) Indirect Illumination

(e) BRDF

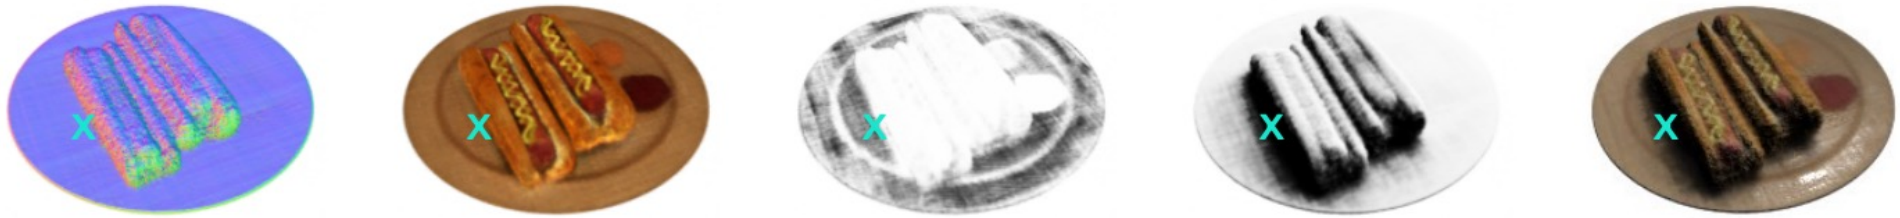

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

(j) Direct

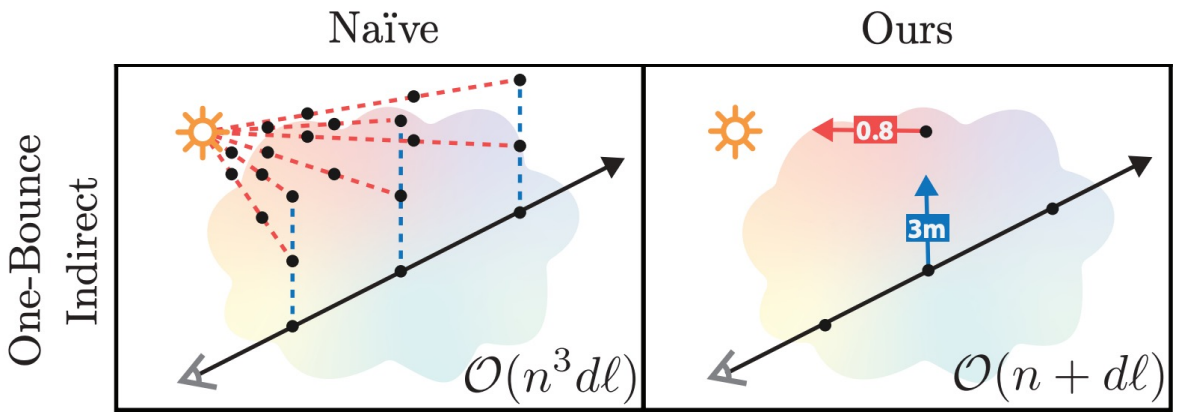

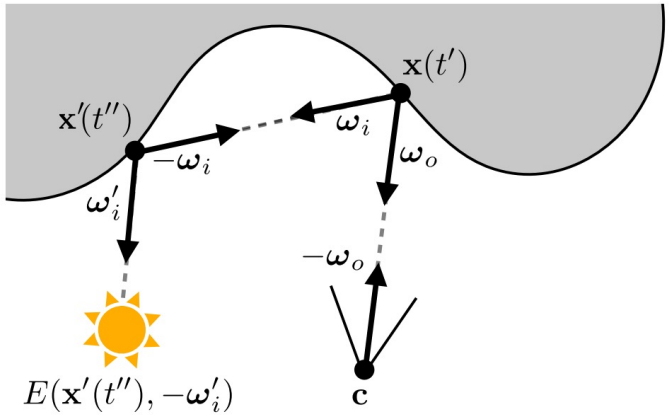

$\mathbf{c}$ = camera center

E = light source

t' = expected termination depth along camera ray

$t'' = \tilde{D}_\phi(\mathbf{x}(t'), \omega_i)$, estimated termination depth along indirect bounce ray

Neural Visibility Fields (NeRV)

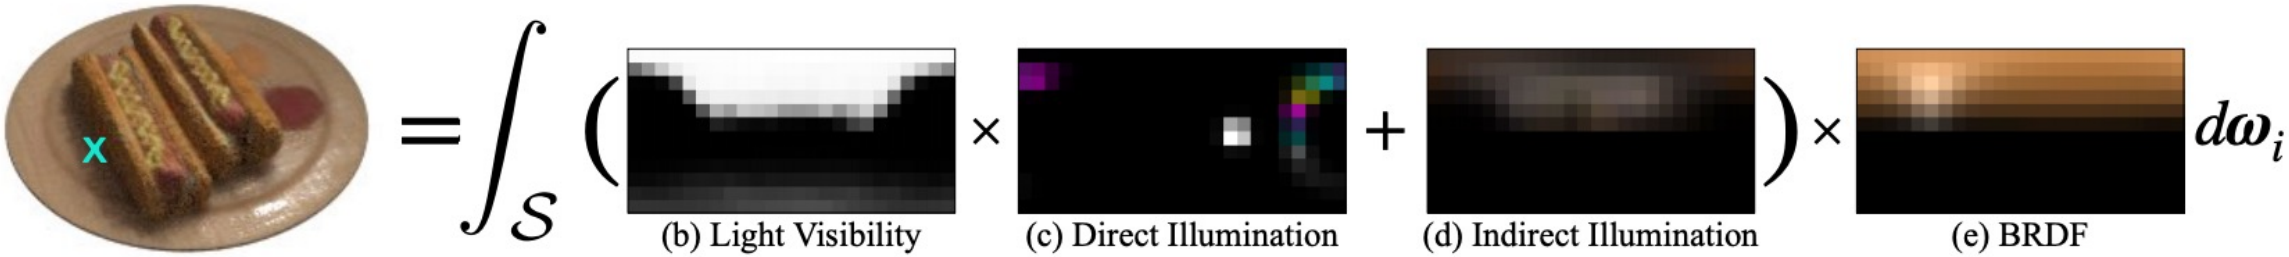

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(d) Indirect Illumination

(e) BRDF

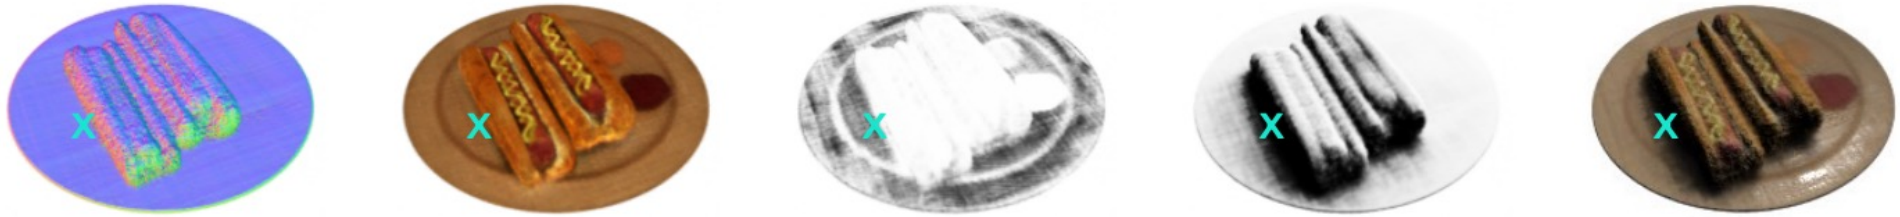

(f) Normals

(g) Albedo

(h) Roughness

(i) Shadow Map

(j) Direct

Neural Visibility Fields (NeRV)

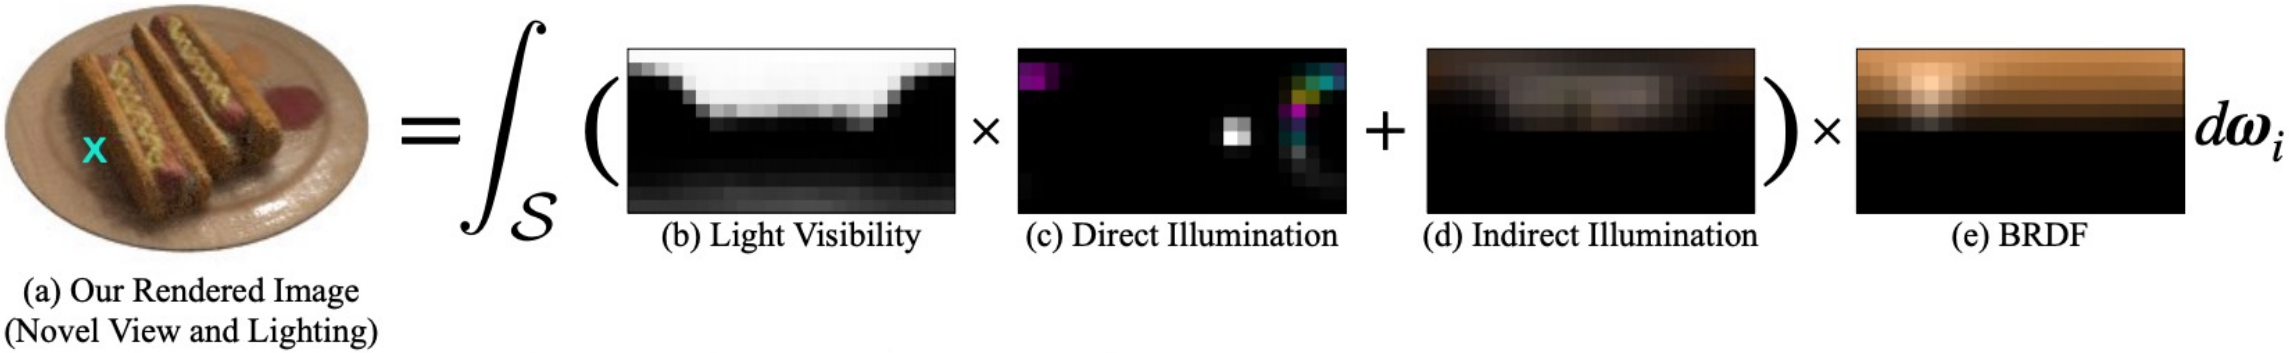

(a) Our Rendered Image (Novel View and Lighting)

(b) Light Visibility

(c) Direct Illumination

(d) Indirect Illumination

(e) BRDF

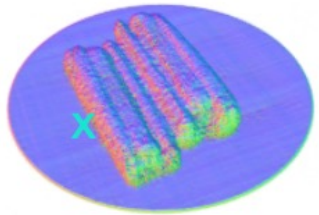

(f) Normals

(g) Albedo

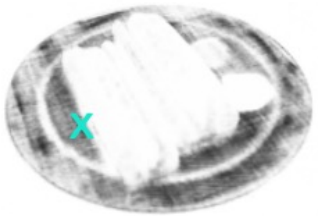

(h) Roughness

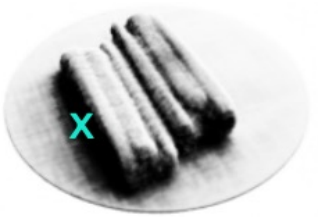

(i) Shadow Map

(j) Direct

(k) Indirect

NeRV models (first-bounce) indirect illumination.

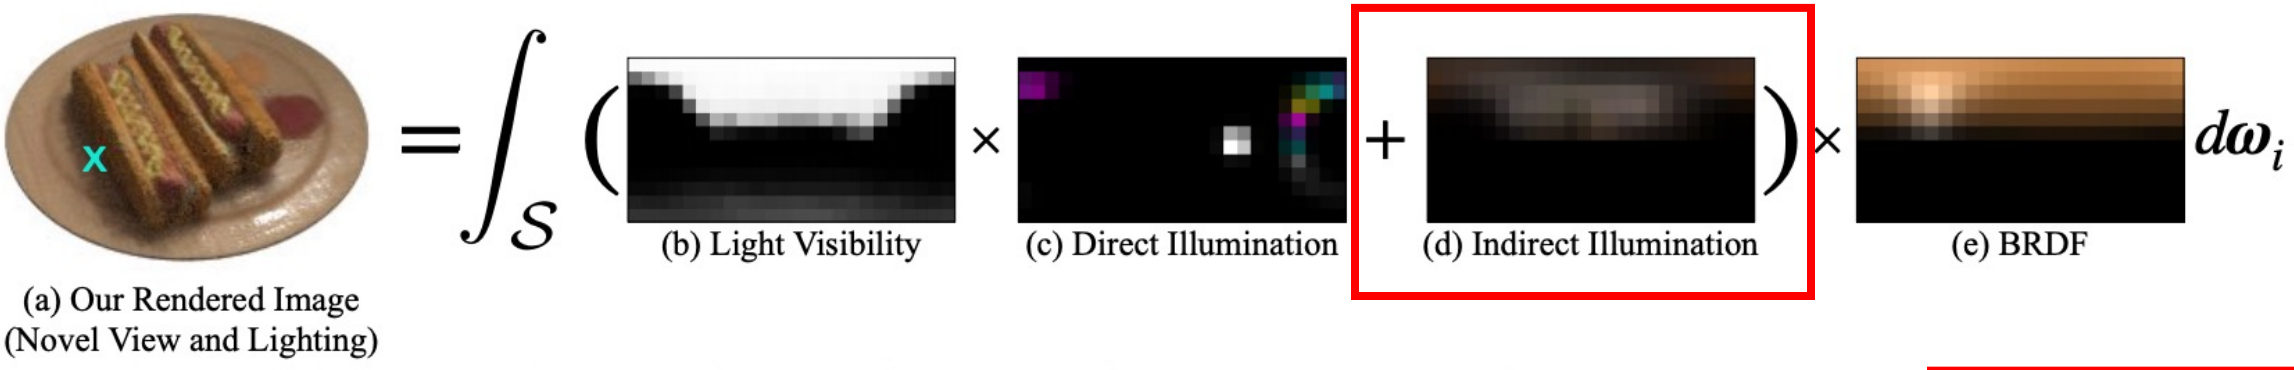

Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, Jonathan T. Barron
NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis
CVPR 2021

NeRV optimizes geometry from scratch.

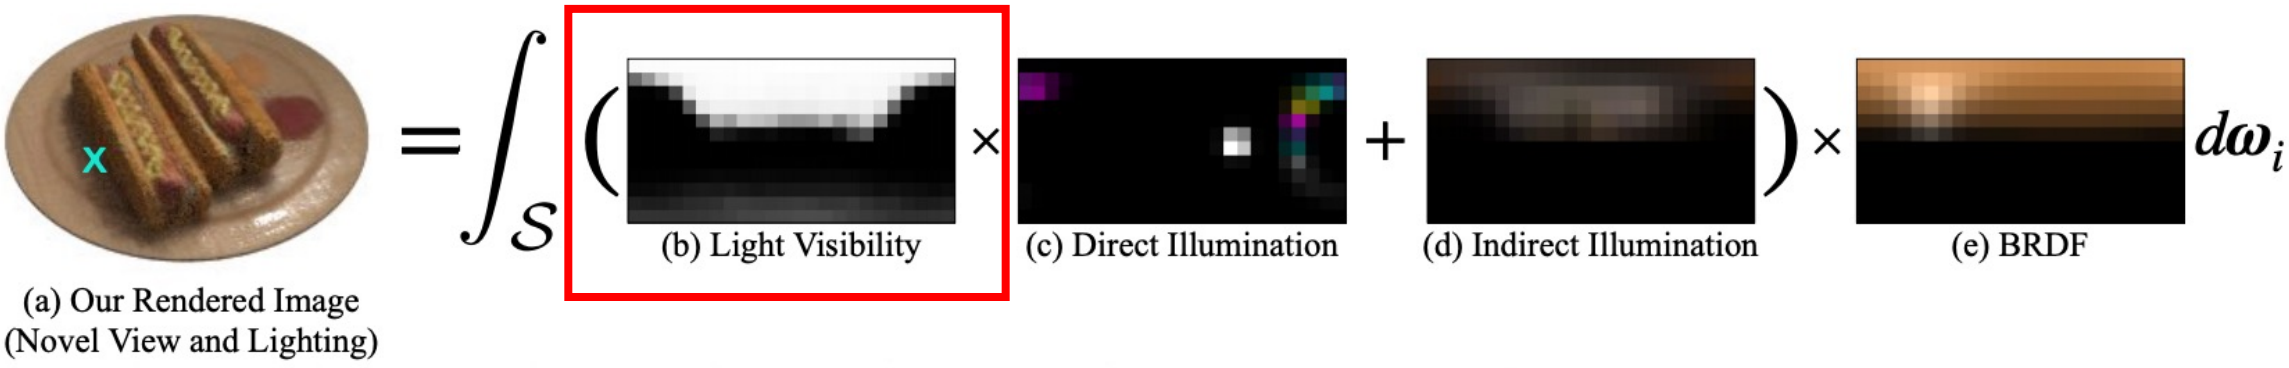

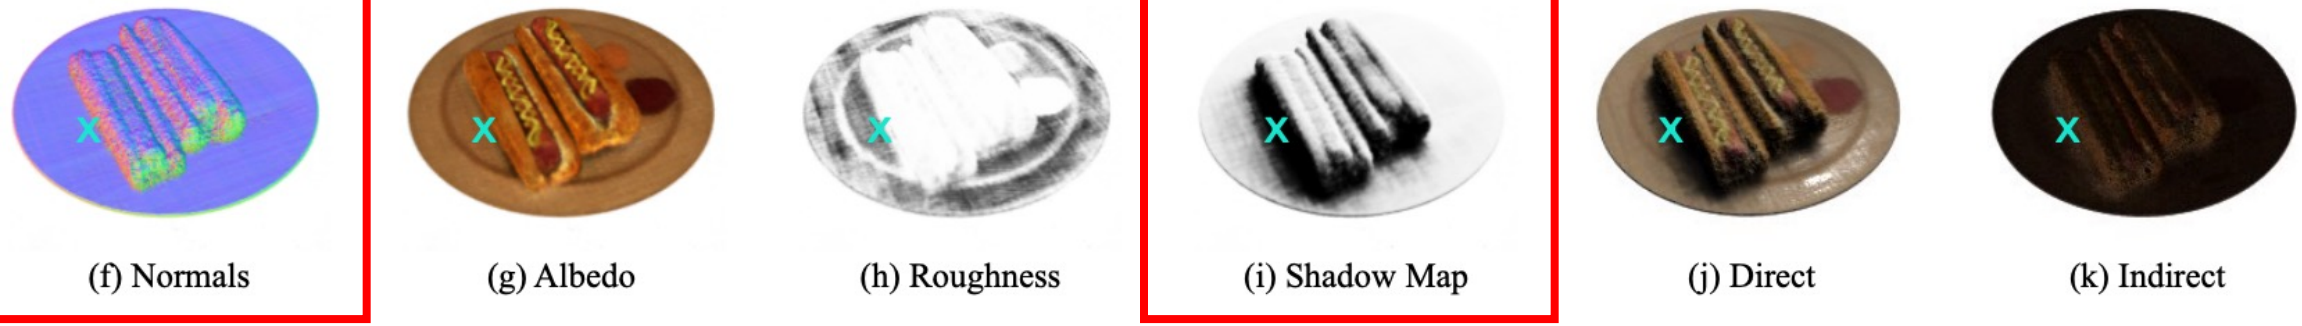

Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, Jonathan T. Barron
NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis
CVPR 2021

NeRV learns from multiple known lighting conditions.

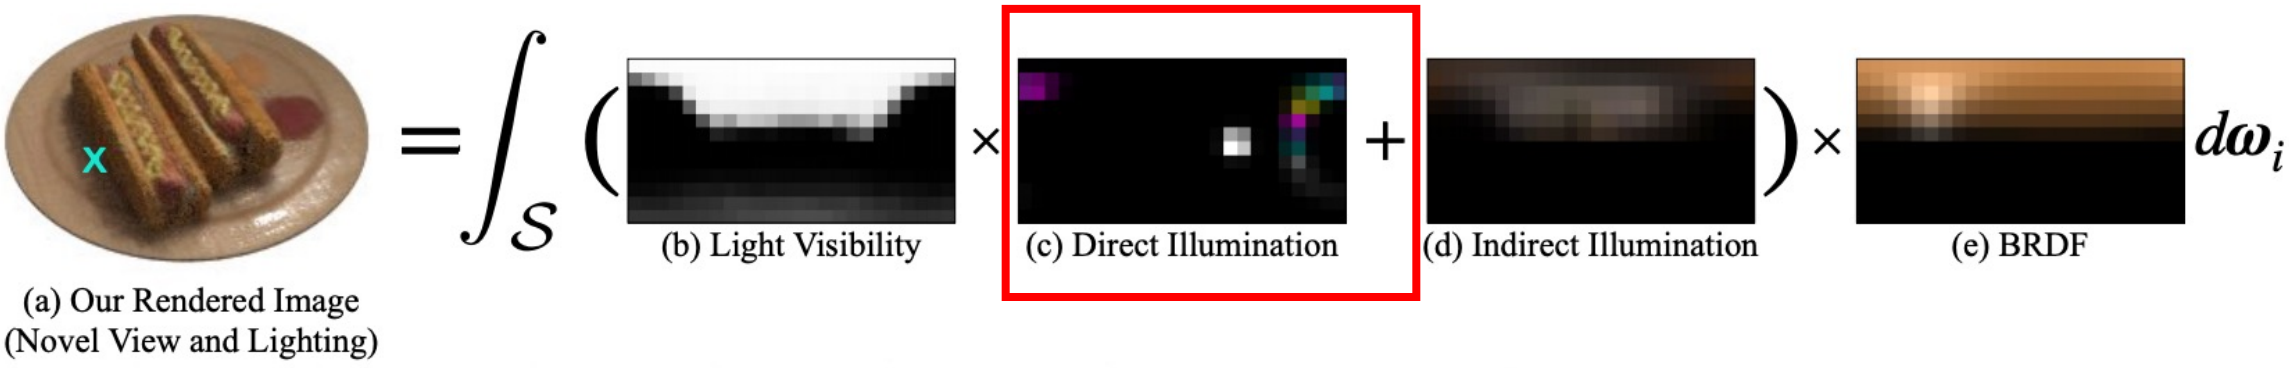

(a) Our Rendered Image
(Novel View and Lighting)

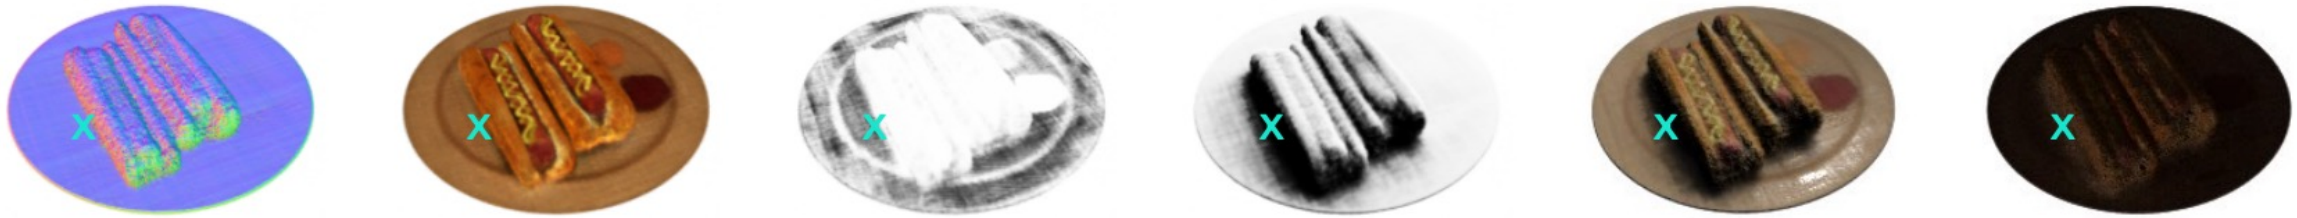

(f) Normals (g) Albedo (h) Roughness (i) Shadow Map (j) Direct (k) Indirect

Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, Jonathan T. Barron
NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis
CVPR 2021

NeRV Results

Novel Views and Lighting



Material Editing


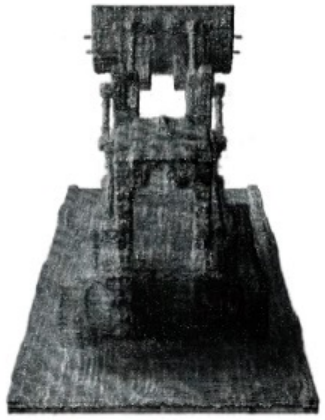


Acknowledgements

Bill Freeman

Jon Barron

Paul Debevec

Pratul Srinivasan

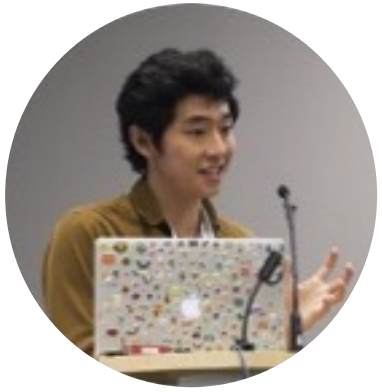

Boyang Deng

Ben Mildenhall

Matt Tancik

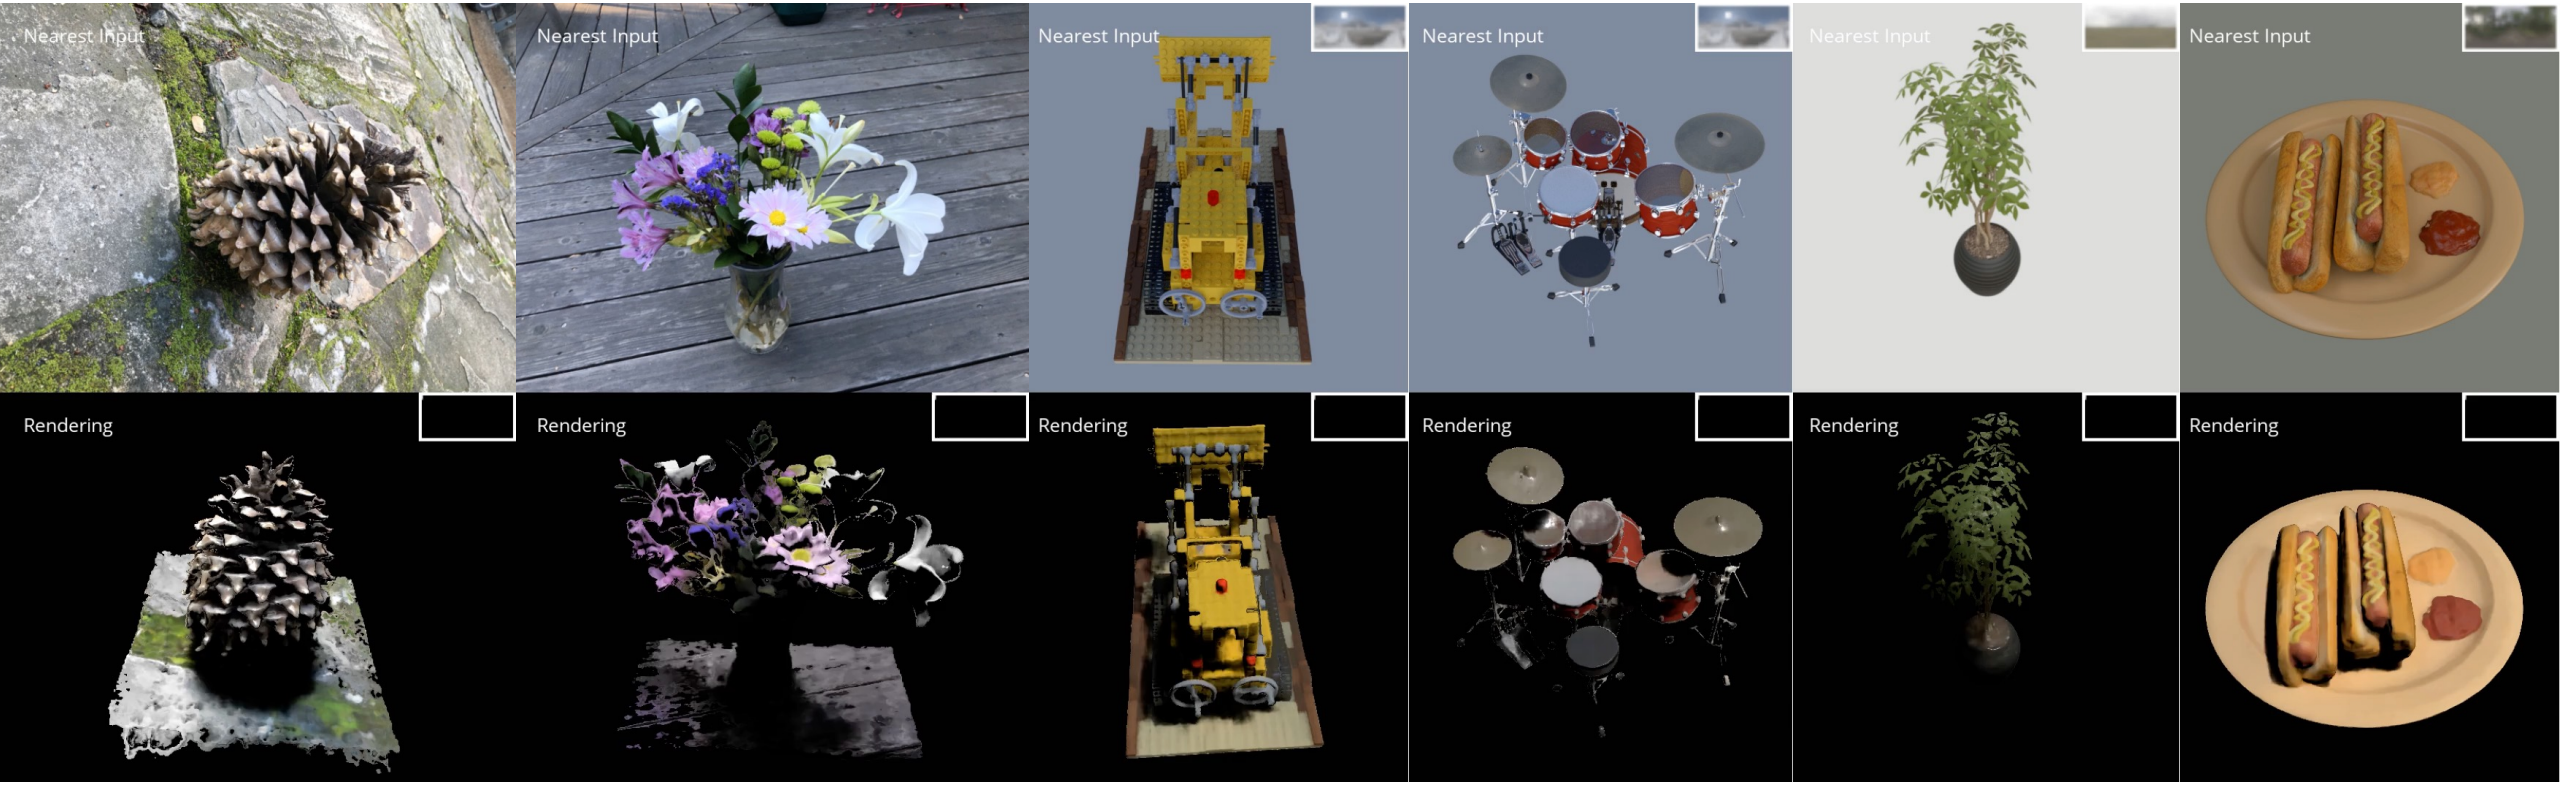

Questions?

NeRFactor:
Paper, videos, code, & data:
https://people.csail.mit.edu/xiuming/projects/nerfactor

Backup

NeRF doesn't require geometry, but doesn't support relighting.

NeRF

Mildenhall*, Srinivasan*, Tancik*, Barron, Ramamoorthi, Ng

ECCV 2020

$$L(\mathbf{c}, \boldsymbol{\omega}_o) = \int_0^\infty V(\mathbf{x}(t), \mathbf{c}) \sigma(\mathbf{x}(t)) L_e(\mathbf{x}(t), \boldsymbol{\omega}_o)\, dt\,,$$

$$V(\mathbf{x}(t), \mathbf{c}) = \exp\left(-\int_0^t \sigma(\mathbf{x}(s))\, ds\right).$$

❌ Not relightable

Neural Reflectance Fields

Bi*, Xu*, Srinivasan, Mildenhall, Sunkavalli, Hašan, Hold-Geoffroy, Kriegman, Ramamoorthi

arXiv 2020

$$L(\mathbf{c}, \boldsymbol{\omega}_o) = \int_0^\infty V(\mathbf{x}(t), \mathbf{c}) \sigma(\mathbf{x}(t)) L_r(\mathbf{x}(t), \boldsymbol{\omega}_o)\, dt\,,$$

$$L_r(\mathbf{x}, \boldsymbol{\omega}_o) = \int_{\mathcal{S}} L_i(\mathbf{x}, \boldsymbol{\omega}_i) R(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)\, d\boldsymbol{\omega}_i\,,$$

❌ Trainable only with lights collocated with cameras

❌ direct illumination only

NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis

Srinivasan, Deng, Zhang, Tancik, Mildenhall, Barron

arXiv 2020

Direct lighting (amortizing visibility computation):

$$\int_0^\infty V(\mathbf{x}(t), \mathbf{c}) \sigma(\mathbf{x}(t)) \int_{\mathcal{S}} \tilde{V}_\phi(\mathbf{x}(t), \boldsymbol{\omega}_i) E(\mathbf{x}(t), -\boldsymbol{\omega}_i) R(\mathbf{x}(t), \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)\, d\boldsymbol{\omega}_i dt$$

Indirect lighting (one bounce only & hard surface assumption):

$$\iint_{\mathcal{S}} \tilde{V}_\phi(\mathbf{x}'(t''), \boldsymbol{\omega}_i') E(\mathbf{x}'(t''), -\boldsymbol{\omega}_i') R(\mathbf{x}'(t''), \boldsymbol{\omega}_i', -\boldsymbol{\omega}_i)\, d\boldsymbol{\omega}_i'$$
$$R(\mathbf{x}(t'), \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)\, d\boldsymbol{\omega}_i$$

✅ Trainable with any lights (including full light probes)

✅ Considers one-bounce indirect illumination

* indicates equal contribution