Representing and Predicting Appearance

Jeppe Revall Frisvad
Technical University of Denmark (DTU)

February 2025

DTU

Appearance representation, fixed lighting

- NeRF: Neural radiance fields
  - Novel view only.
  - Ray marching to get points.

[Mildenhall *et al.* 2020. ECCV]

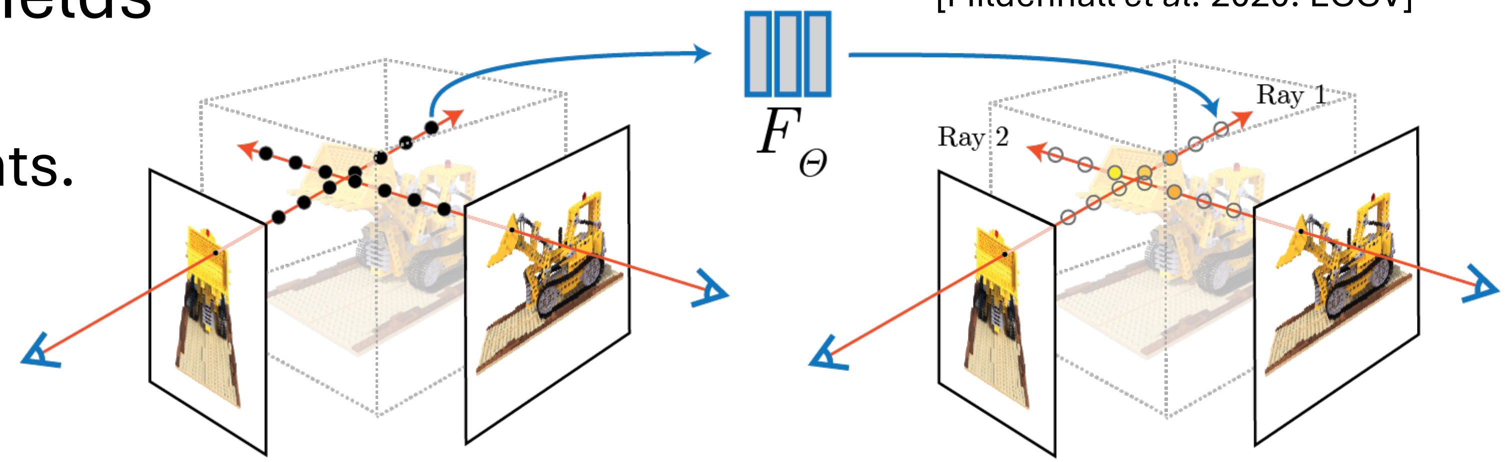

- Gaussian splatting
  - Novel view only.
  - Adaptive positioning of points along surfaces.

[Kerbl *et al.* 2023. SIGGRAPH]

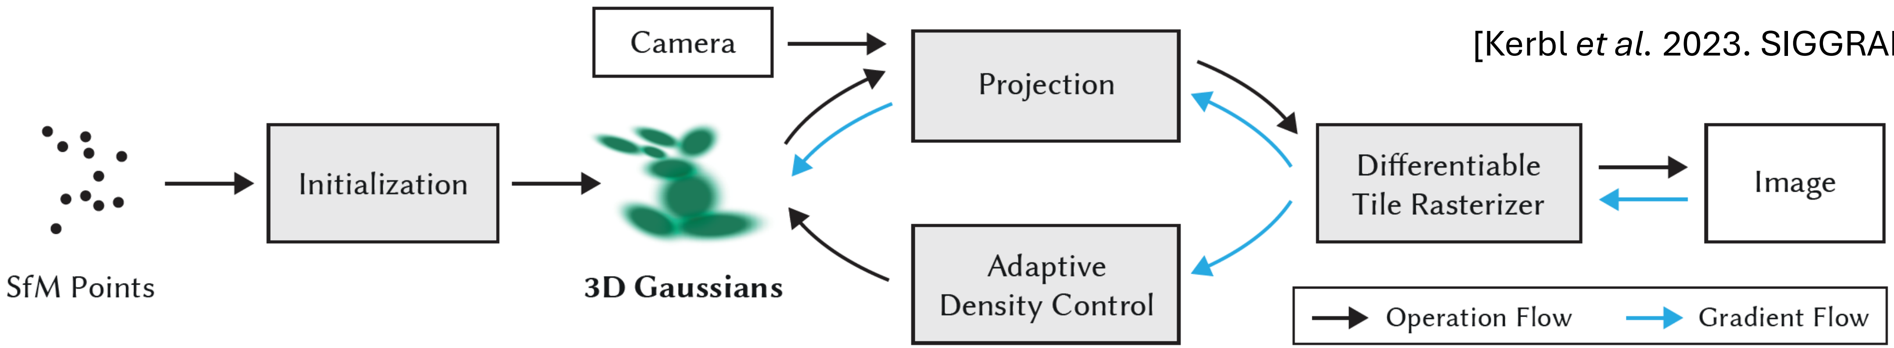

Appearance representation, relightable

- NeRFactor
  - BRDF assumption.
  - No translucency.

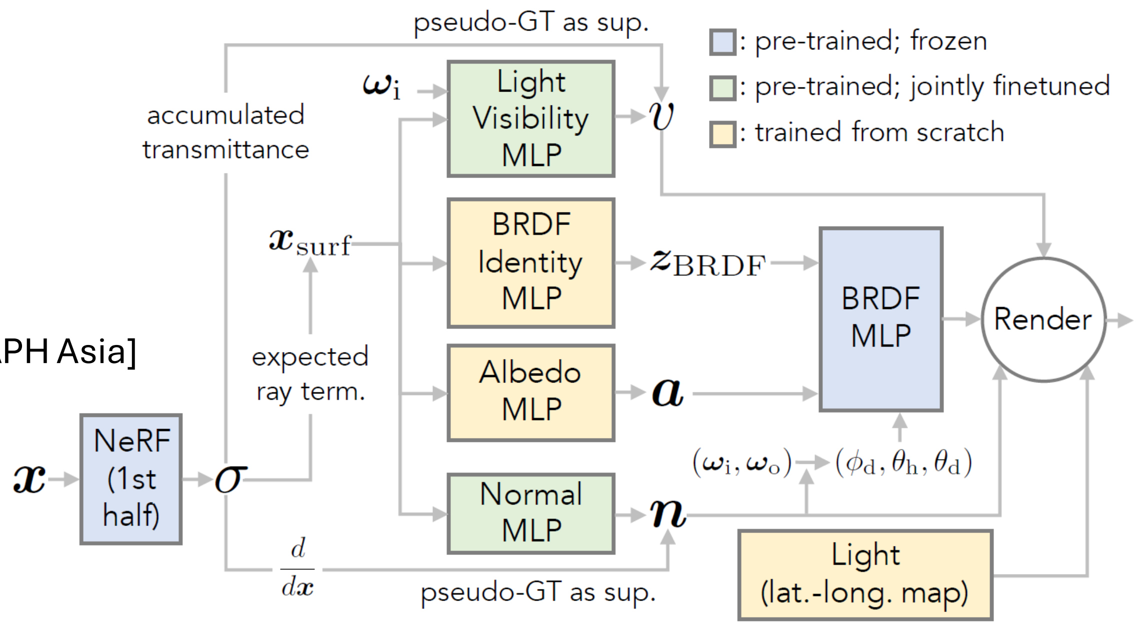

[Zhang *et al*. 2021. SIGGRAPH Asia]

- NeuMIP
  - Texture mapping required.
  - Directional lights only.
  - Flat patch translucency.

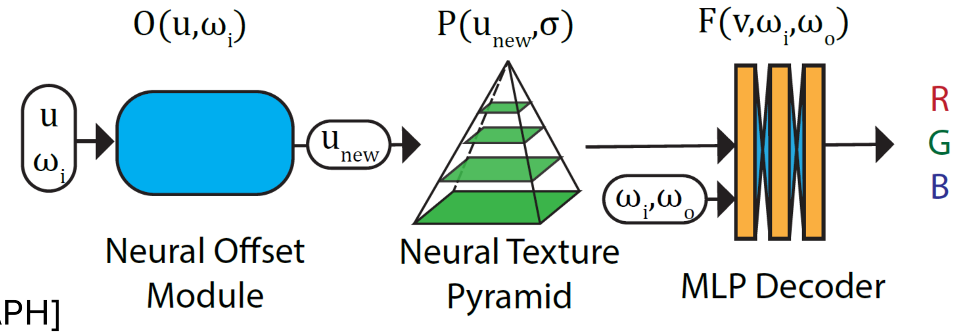

[Kuznetsov *et al*. 2021. SIGGRAPH]

Appearance representation, translucency

- Textured diffuse dipole BSSRDF
  - Surface variation assumption.
  - Non-directional translucency.

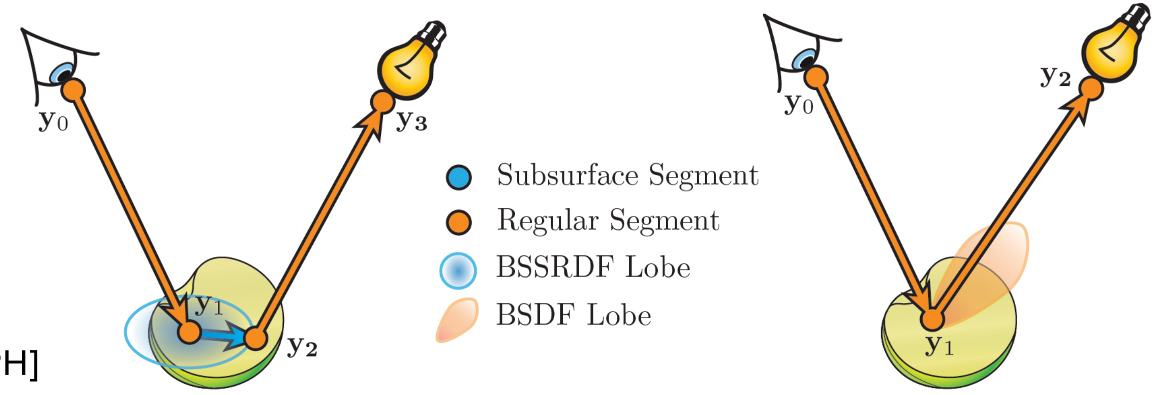

[Deng *et al*. 2022. SIGGRAPH]

- NeuralTO: Neural translucent objects
  - Better geometry.
  - Constant extinction.
  - No refraction.
  - Scatter simplifications.
  - Relightable?

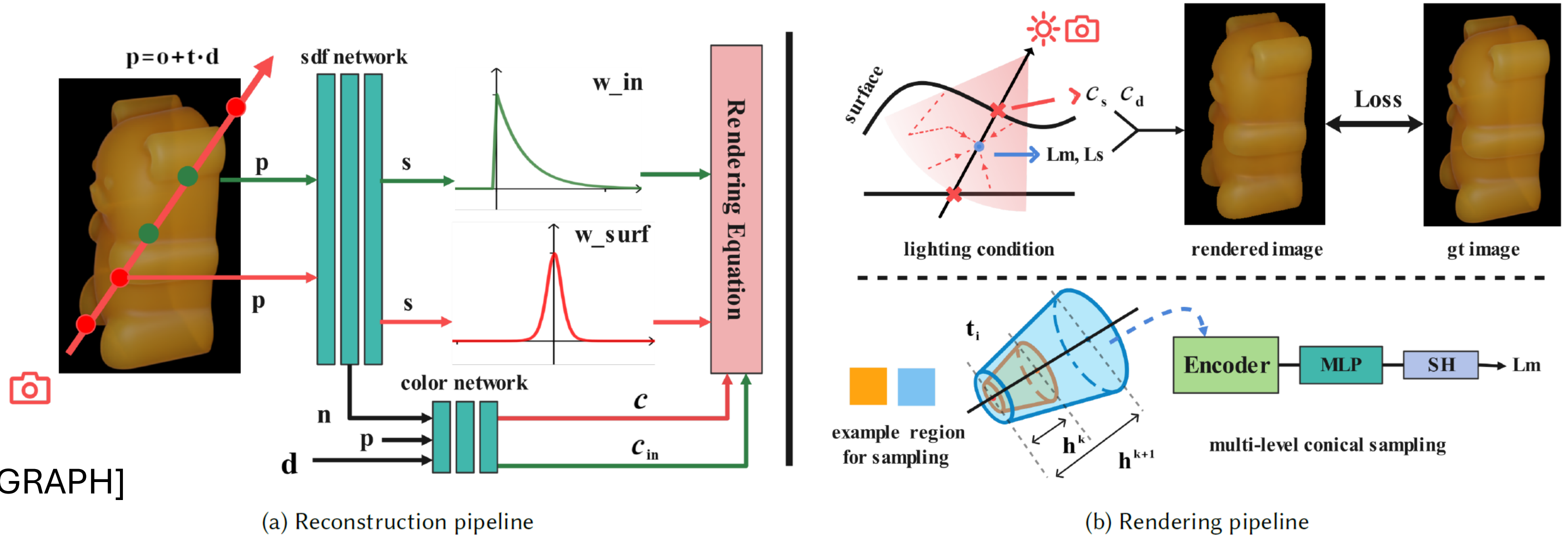

(a) Reconstruction pipeline

(b) Rendering pipeline

[Cai *et al*. 2024. SIGGRAPH]

Appearance representation, simple lighting

- Radiance predicting neural networks
  - Directional lighting, no refraction, per scattering event, or
  - Diffuse lighting.

[Kallweit *et al.* 2017. SIGGRAPH Asia]

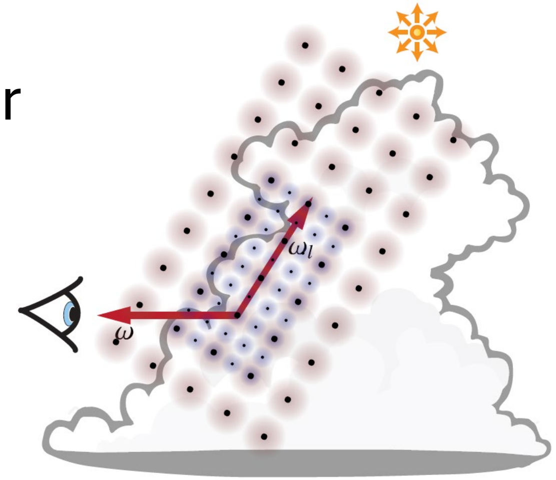

[Rittig *et al.* 2021. EG]

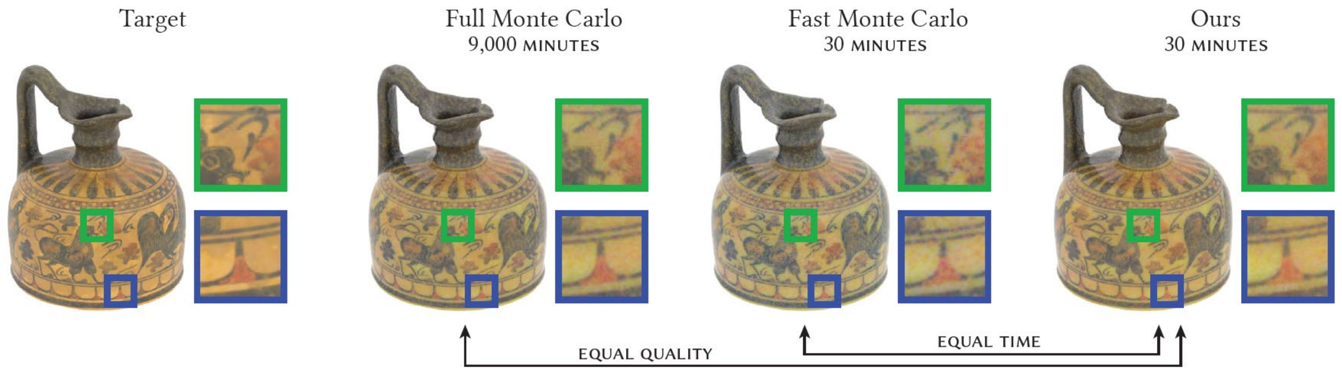

- Relightable NeRF
  - Point lighting.
  - No refraction.

[Zeng *et al.* 2023. SIGGRAPH]

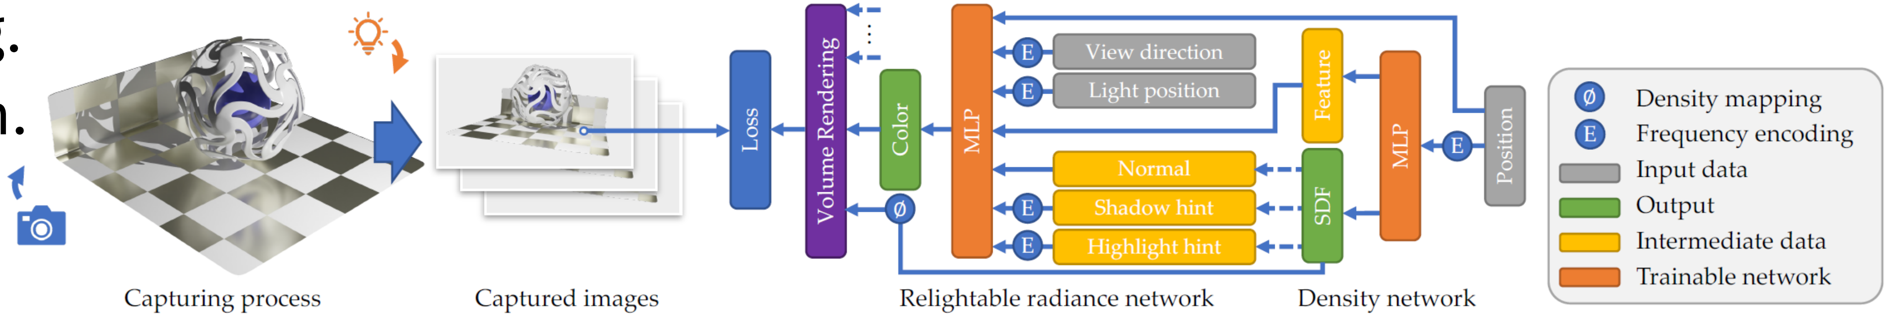

Thomson TG

- **Research project**
  PRIME: Predictive Rendering in Manufacture and Engineering
- **PhD project**
  Macroscopic Appearance Specification and Rendering
- **Papers** (included in the following)

NeuPreSS: compact neural precomputed subsurface scattering for distant lighting of heterogeneous translucent objects
Thomson TG, Jeppe Revall Frisvad, Ravi Ramamoorthi, Henrik Wann Jensen
Computer Graphics Forum (PG 2024) 43(7), Article e15234. October 2024.

Neural SSS: lightweight object appearance representation
Thomson TG, Duc Minh Tran, Henrik Wann Jensen,
Ravi Ramamoorthi, Jeppe Revall Frisvad
Computer Graphics Forum (EGSR 2024) 43(4), Article e15158.
July 2024.

Appearance representation, known geometry

- NeuPreSS: Neural Precomputed Subsurface Scattering
  - Directional lighting.
  - Known geometry.
  - Separate surface reflection.
  - Conversion to SH-based PRT.
  - Expensive training.
  - Learn to importance sample.

[TG *et al.* 2024. PG]

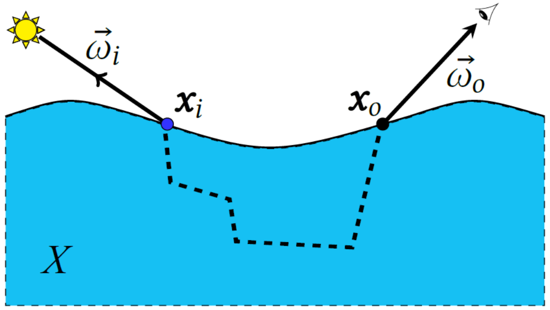

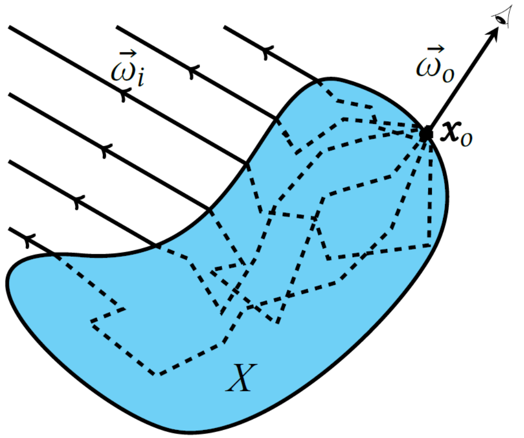

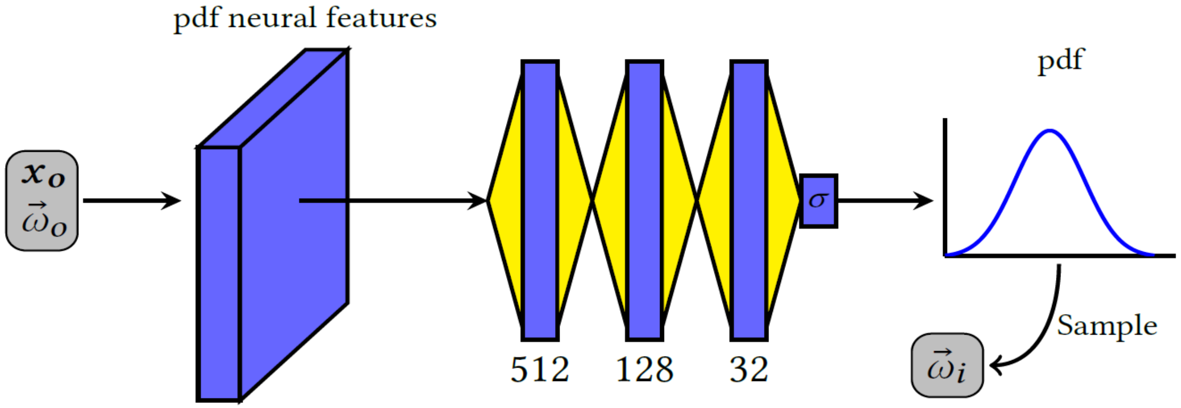

(a) Importance Sampling Module

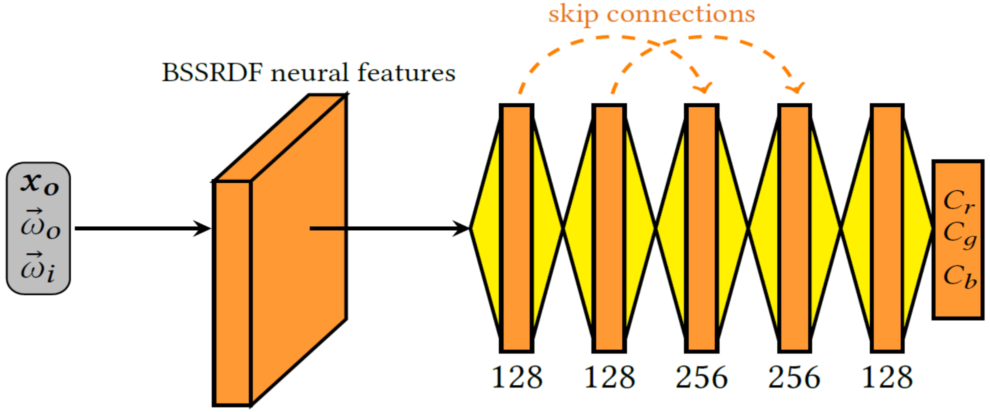

(b) Appearance Specification Module

Comparison, multi-sampled directional lighting

- Representing the appearance of a digital object.

- References: path tracing of a heterogeneous volume with a refractive interface.

- Numbers are rendering times for 1 sample per pixel.

- Images (except references) were rendered using 256 samples per pixel.

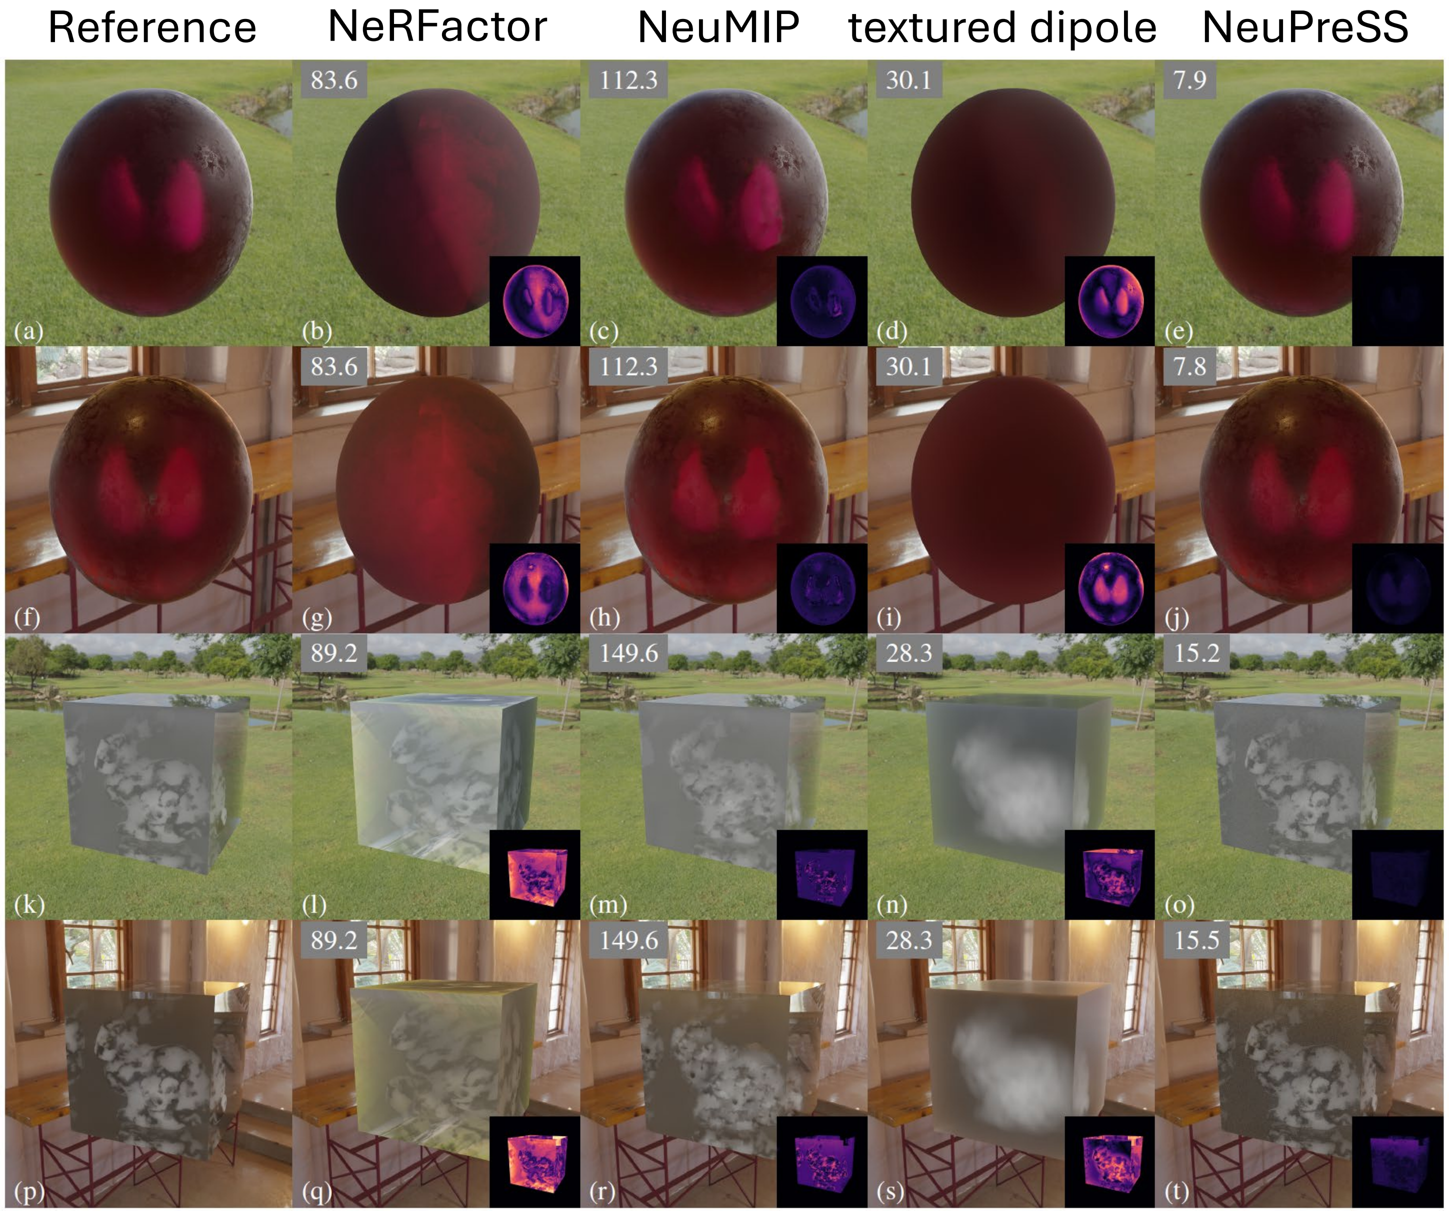

[TG *et al*. 2024. PG]

Appearance representation, known geometry

- Neural SSS
  - Known geometry.
  - Separate surface reflection.
  - Inexpensive training.

[TG *et al.* 2024. EGSR]

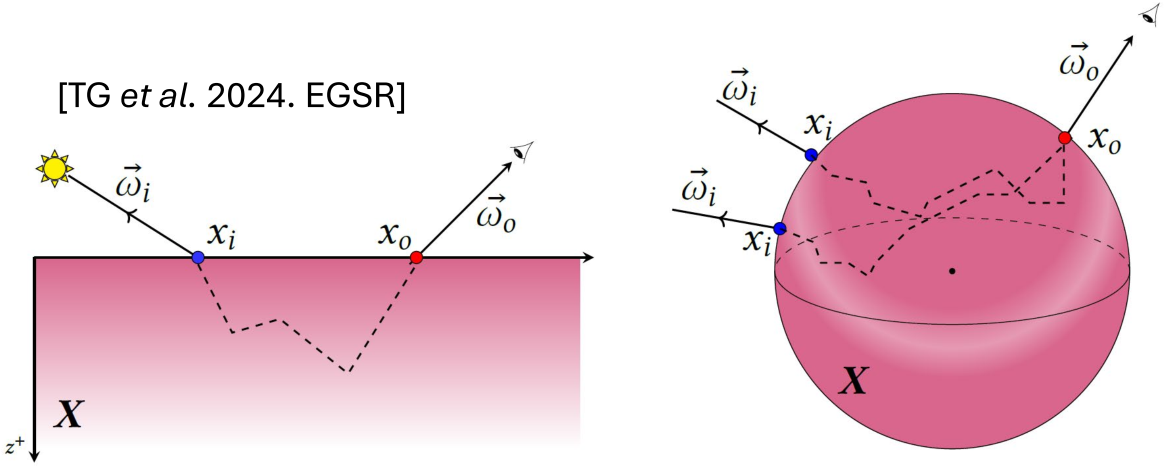

- Trained using non-converged unidirectional volume path tracing!
  - Train while rendering.
  - Switch to N-samples neural BSSRDF when trained.
  - Normalizing flow using network and scene for importance sampling.

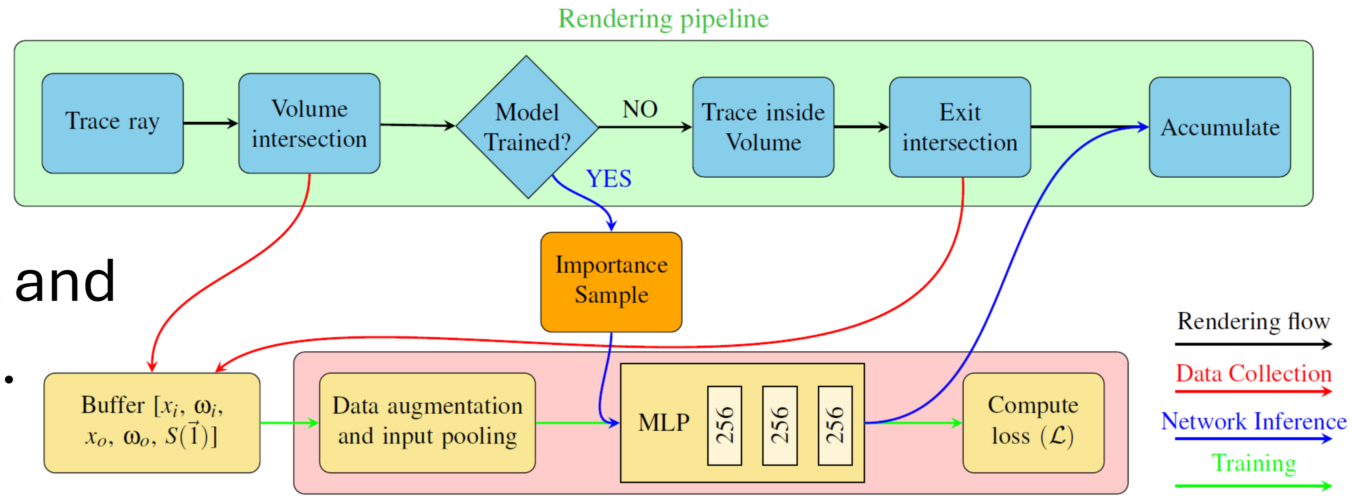

$$L_r(\boldsymbol{x}_o, \vec{\omega}_o) = \int_G \int_{4\pi} \left(\boldsymbol{S}_N \otimes \boldsymbol{L}_{i,N}\right)^T \mathrm{diag}\left(\boldsymbol{n}_{i,N}\, \boldsymbol{\omega}_{i,N}^T\right) \mathrm{d}\omega_i \, \mathrm{d}A_i$$

Comparison, global illumination

- Representing the appearance of a digital object.

[TG *et al*. 2024. EGSR]

- References: path tracing of a heterogeneous volume with a refractive interface.

- Deviations due to surface texture mapping (textured dipole) and distant lighting (NeuPreSS) are as expected.

Reference | textured dipole | NeuPreSS | Neural SSS

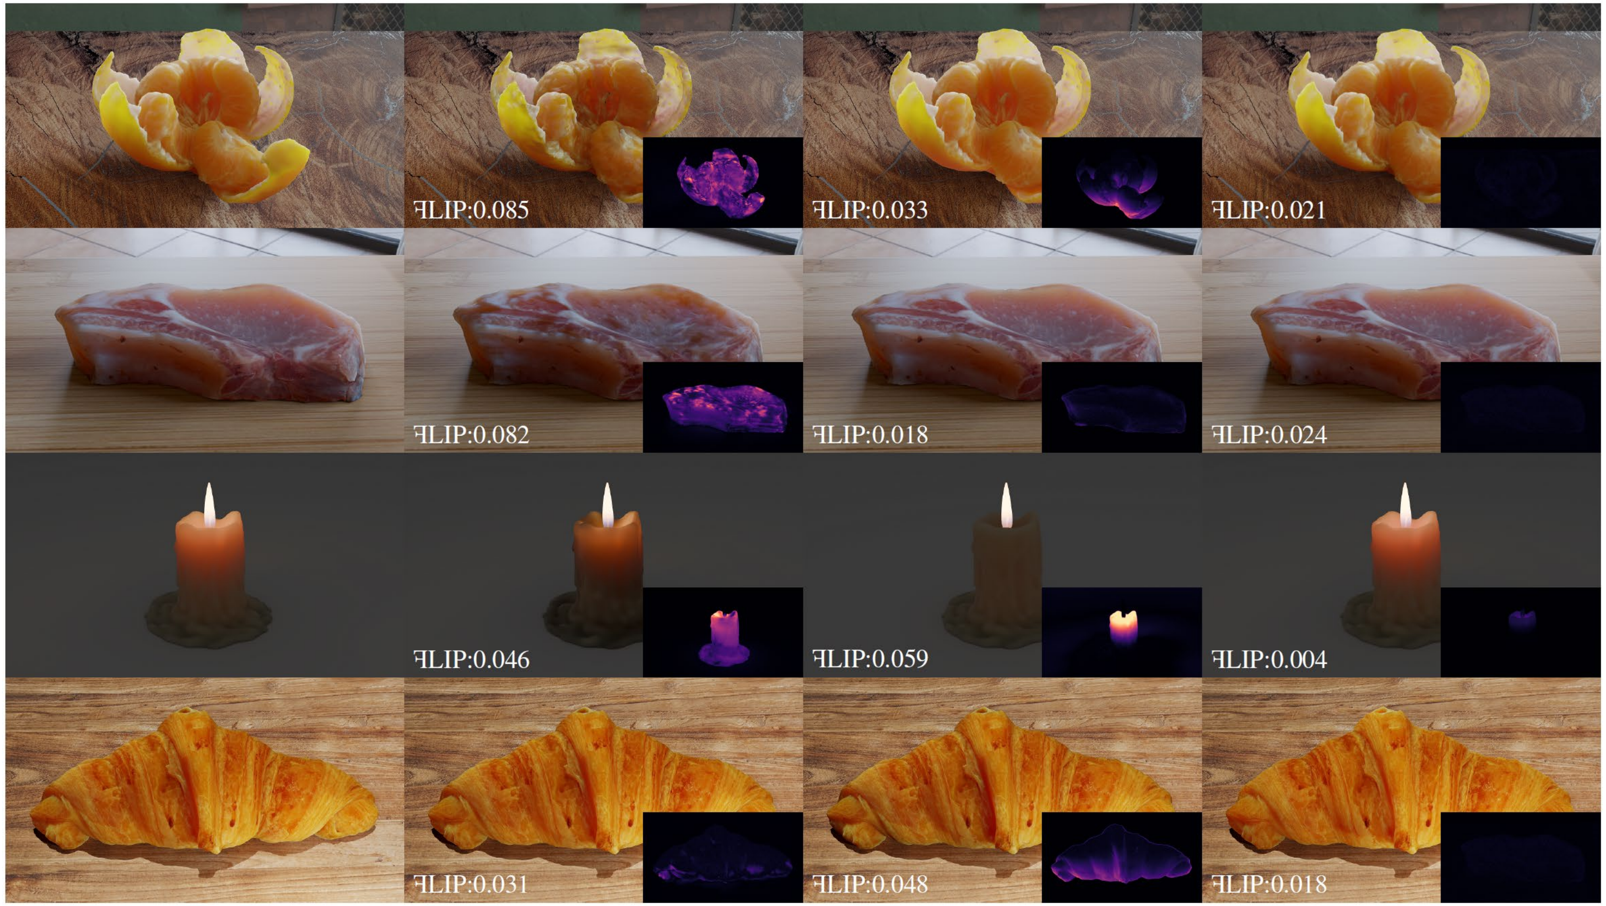

Material appearance prediction

- How to predict the appearance of an unseen object?
- Physically based rendering is good but how good?
- Validate by modeling digital scenes that match physical scenes?

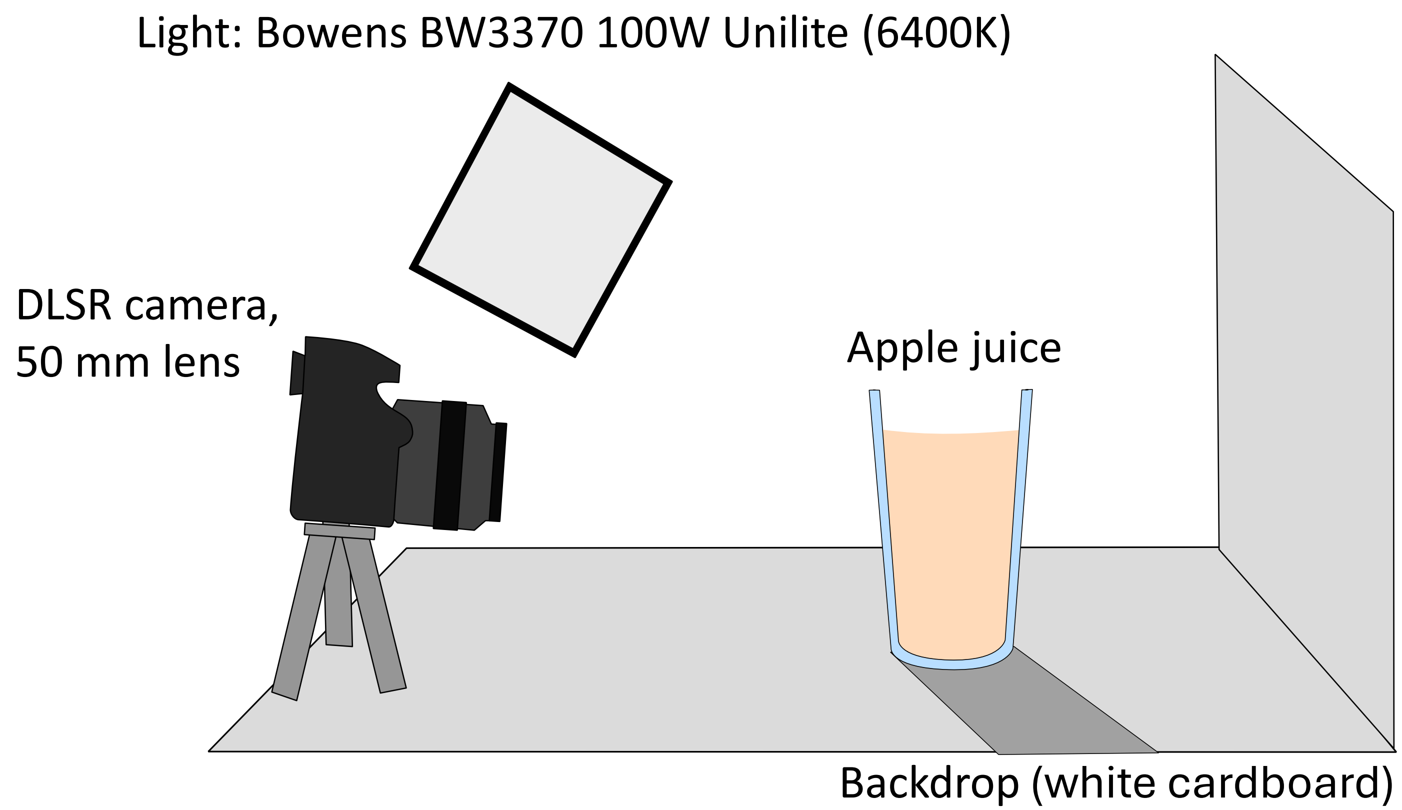

photograph | CAD model rendering | absolute difference × 2

[Dal Corso *et al.* 2016. MAM]

Appearance prediction, editability

water vitamin B2 protein fat skimmed low fat whole

- Digitizing material appearance: accurate intrinsic optical properties required.
- Industry standard: plausible appearance for entertainment.
- Industry need: predictive appearance of a manufactured item (visualizing the digital twin).
- Research challenge: **editable digital representations of real objects**.
- Important aspects: **modeling** (math and physics), **validation** (measurements), **acquisition** (vision and inverse methods), **application** (quality control, prototyping, etc.).

Multiscale modelling

[Frisvad *et al.* 2020. EG]

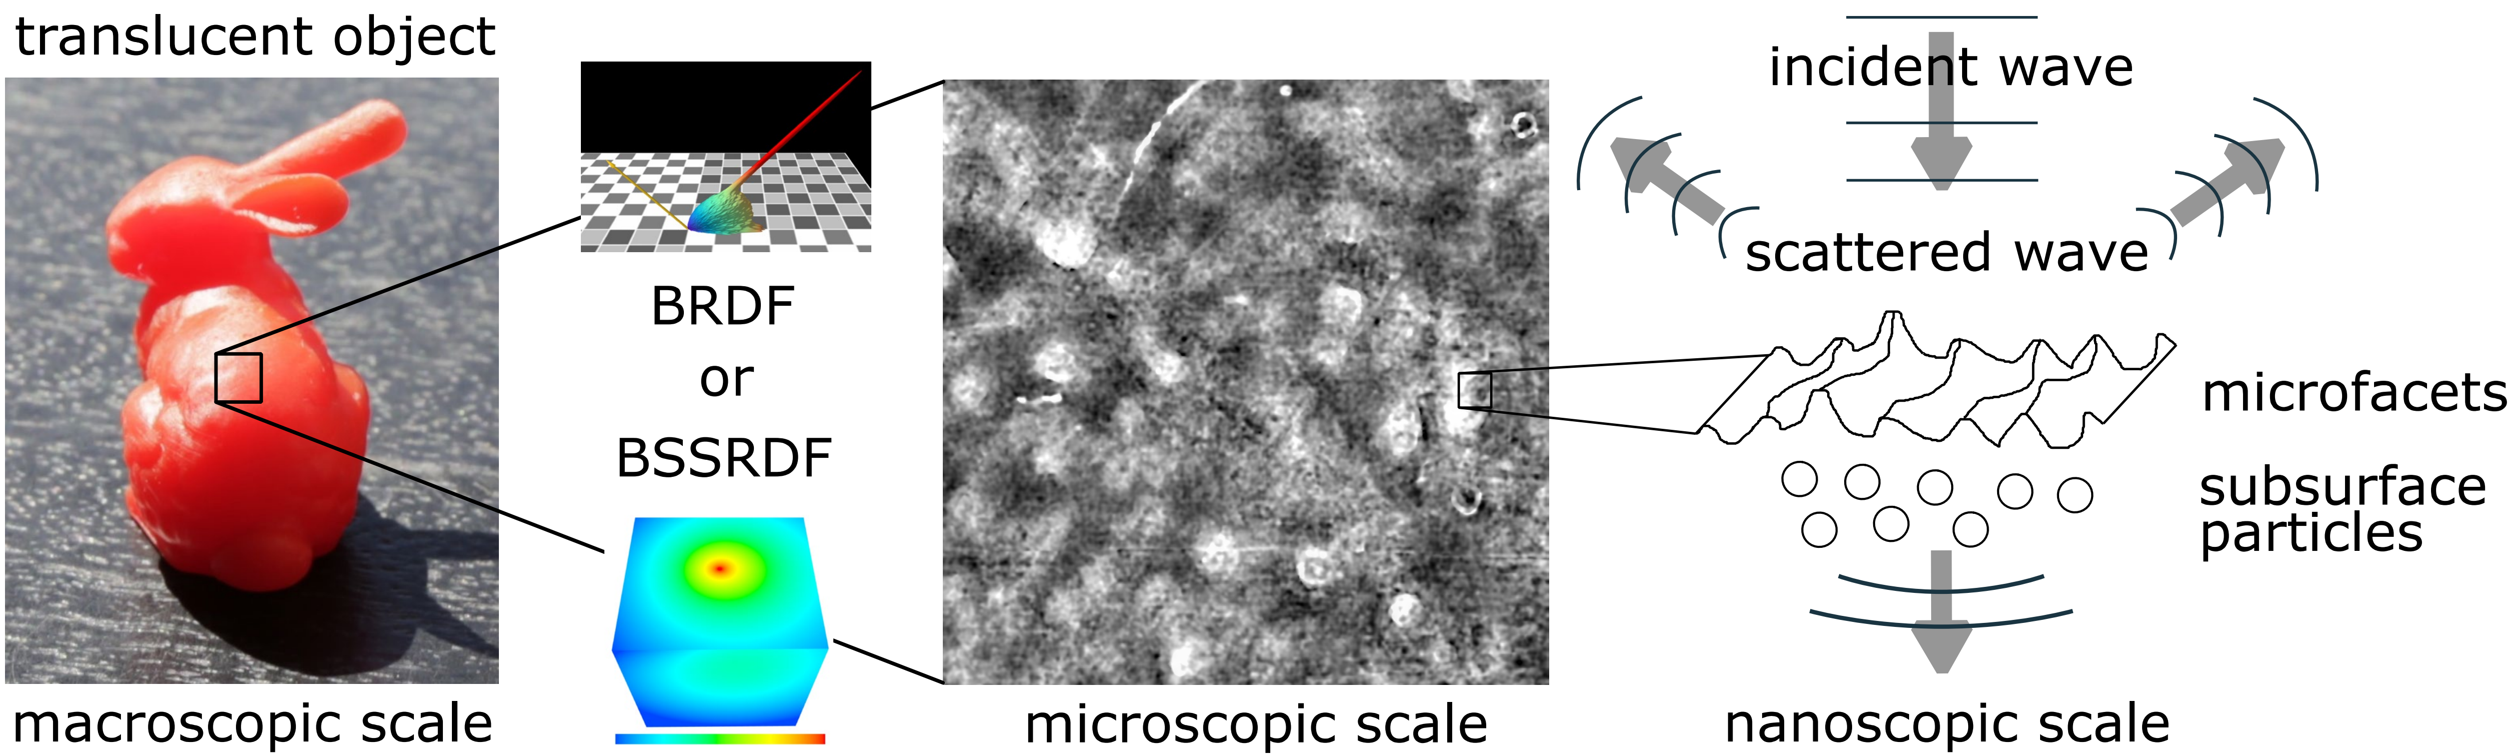

- With simulation of light propagation, we can compute **macroscopic optical properties** by considering geometry at different scales.

Models at different scales

- Microscopic scale:
  - Nano/micro: models considering explicit microgeometry.
  - Micro/milli: models using particle size or microfacet normal distribution functions.
- Macroscopic scale:
  - BSSRDF: models where the points of incidence and emergence are different.
  - BRDF/BTDF: local models for opaque/thin objects.

microsurface

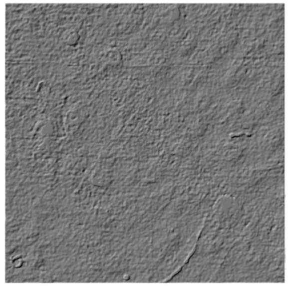

profilometry

particles

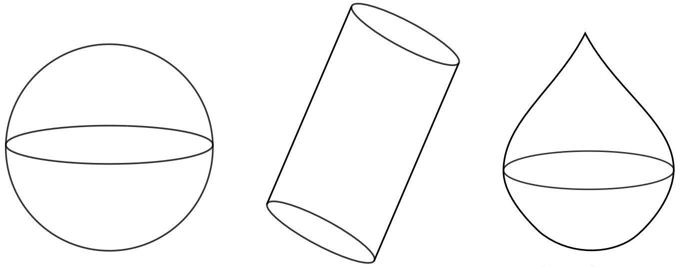

sphere cylinder raindrop

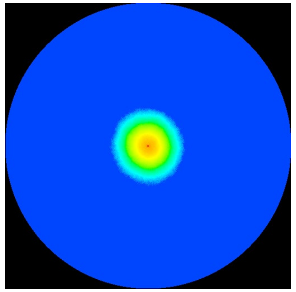

normal distribution

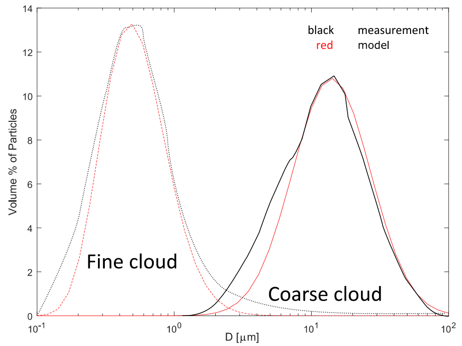

particle size distribution

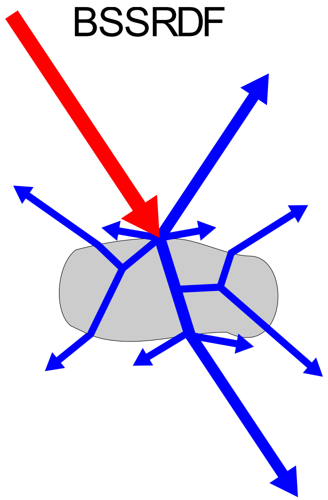

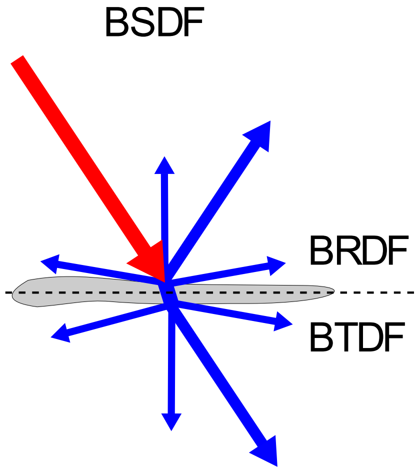

Index of refraction (or refractive index)

- Combining permittivity (ε), permeability (μ), and conductivity (σ):
- $n_{\text{med}} = n' + i\,n'' = c\sqrt{\mu\left(\varepsilon + i\,\frac{\sigma}{\omega}\right)}$
    - ω is angular frequency.
    - c is the speed of light *in vacuo*.
- Real part $n' \approx \frac{c}{v}$
    - v is the phase velocity of the light wave.
- Imaginary part $n'' \approx \frac{\sigma_a \lambda}{4\pi}$
    - σ_a is the absorption coefficient.
    - λ is the wavelength *in vacuo*.

varying the real part n'

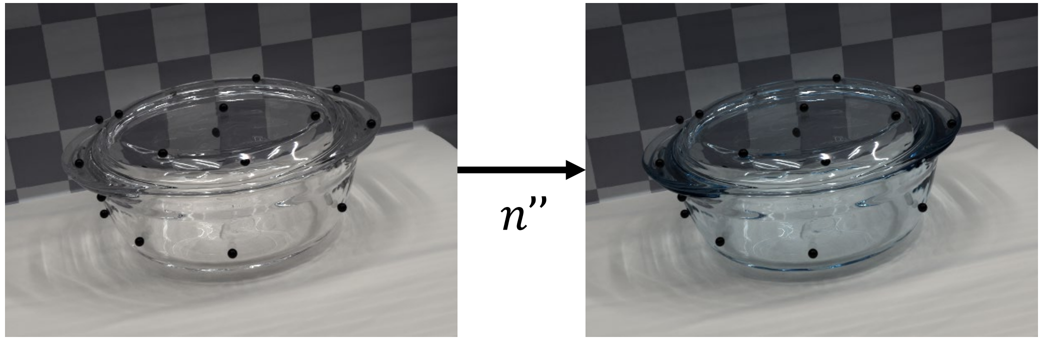

Including absorption

[Stets *et al.* 2017. Applied Optics]

Microfacet BSDF

- A surface is **optically smooth** if the surface roughness R_q is sufficiently small compared with the wavelength λ.
- Rayleigh smooth-surface criterion: $R_q < \lambda/(8\cos\theta_i)$.
- Considering smooth microgeometry we can use n_{med} as input for analytic or computational solutions for Maxwell's equations.
- Example: Fresnel reflectance F for a microfacet BSDF.

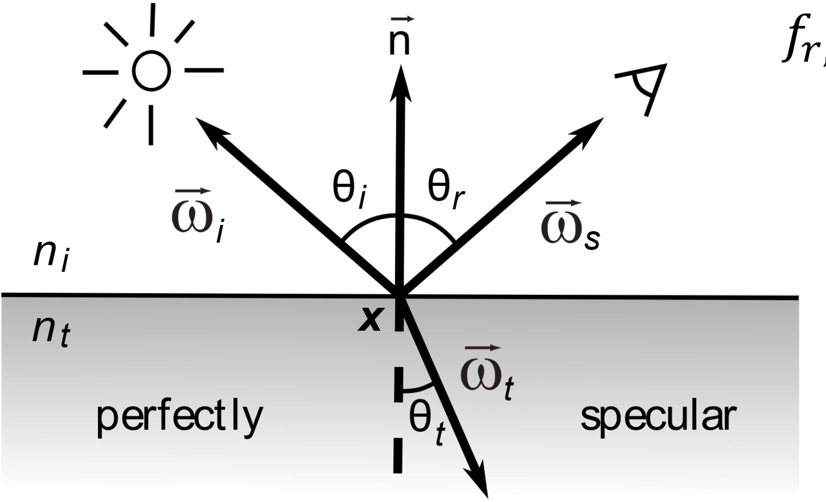

$$f_{r,m}(\boldsymbol{x}_m, \vec{\omega}_i, \vec{\omega}_o) = F\left(\vec{m}, \vec{\omega}_i, \frac{n_t}{n_i}\right) \frac{\delta(\vec{\omega}_o - \vec{\omega}_s)}{\cos\theta_i}$$

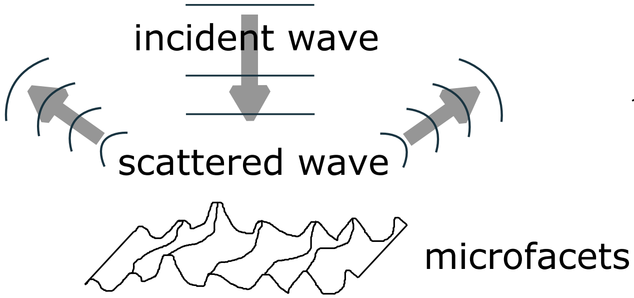

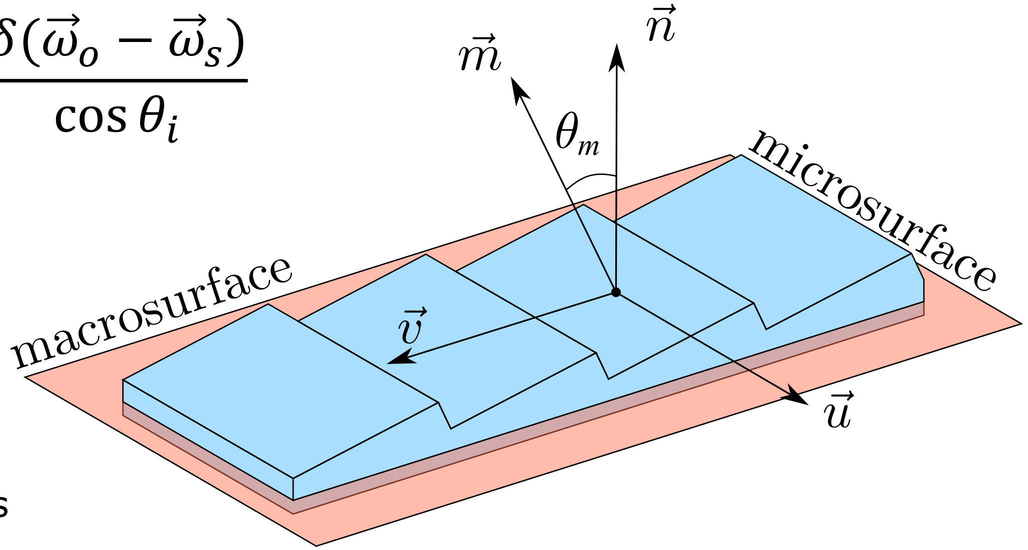

Particle phase function and cross sections

- Particle cross sections
  - C_g is the geometric cross section.
  - C_s is the scattering cross section.
  - C_a is the absorption cross section.
  - $C_t = C_s + C_a$ is the extinction cross section.
- Particle phase function
  - $p_m(\vec{\omega}_i, \vec{\omega}_o)$ is the far field distribution of the scattered light.
  - $g = \int_{4\pi} p_m(\vec{\omega}_i, \vec{\omega}_o)\,(\vec{\omega}_i \cdot \vec{\omega}_o)\,\mathrm{d}\omega$ is the asymmetry parameter in [-1,1].

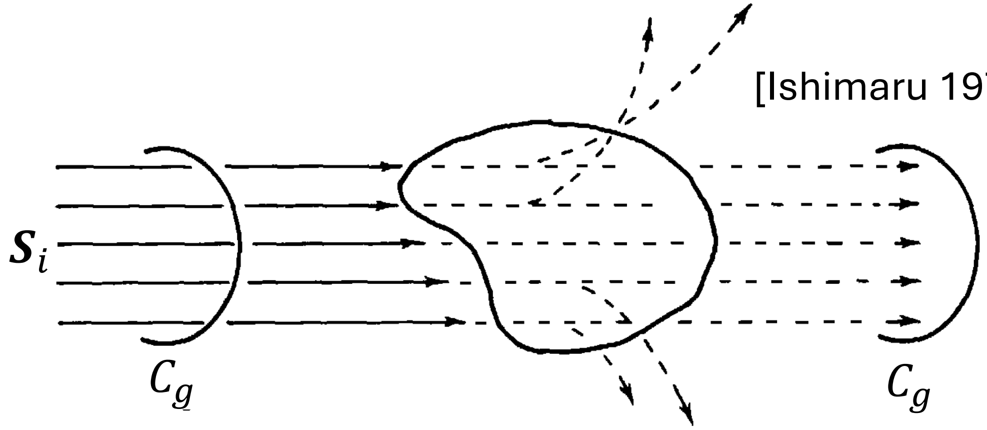

[Ishimaru 1978]

Example: Insert $x = \frac{2\pi r n_{\text{med}}}{\lambda}$ and $y = \frac{2\pi r n_p}{\lambda}$ in Lorenz-Mie theory to compute C_s, C_t, and p of a spherical particle of radius r.

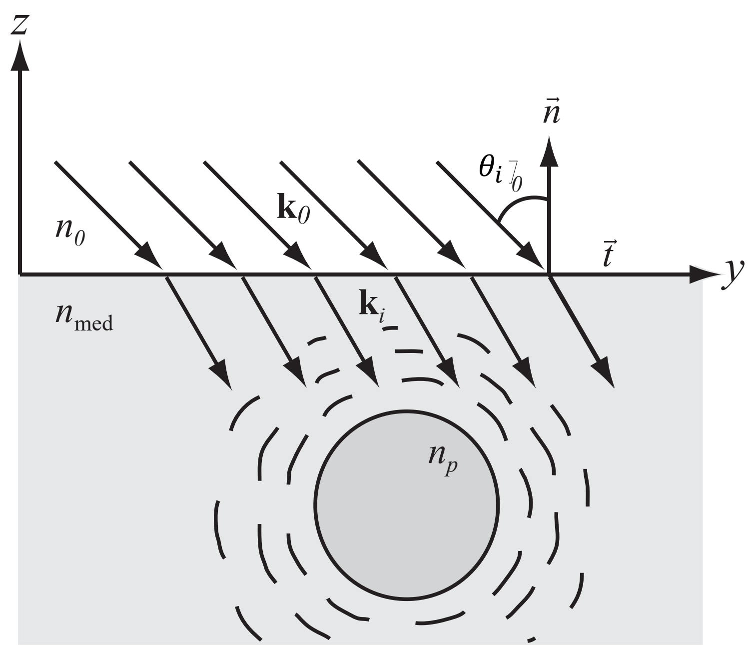

small particle

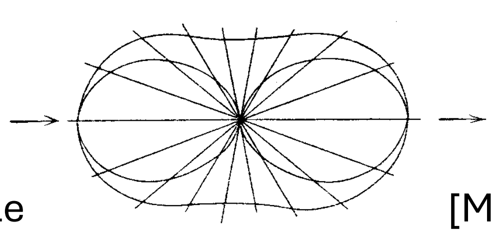

[Mie 1908]

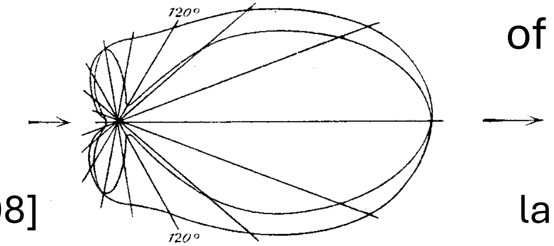

large particle

Scattering properties of a medium

[Frisvad *et al*. 2007. SIGGRAPH]

- Using a particle size distribution $N(r)$: $\sigma_s = \int_{r_{\min}}^{r_{\max}} C_s(r) N(r)\,\mathrm{d}r$
  - σ_s is the scattering coefficient.
  - Similarly for σ_a (absorption coefficient) and p (ensemble phase function).
- Using a microfacet normal distribution $D(\vec{m})$:

[Walter *et al*. 2007. EGSR]

$$f_s(\vec{\omega}_i, \vec{\omega}_o, \vec{n}\,) = \int \left|\frac{\vec{\omega}_i \cdot \vec{m}}{\vec{\omega}_i \cdot \vec{n}}\right| f_m(\vec{\omega}_i, \vec{\omega}_o, \vec{m}) \left|\frac{\vec{\omega}_o \cdot \vec{m}}{\vec{\omega}_o \cdot \vec{n}}\right| G(\vec{\omega}_i, \vec{\omega}_o, \vec{m})\, D(\vec{m})\, \mathrm{d}\omega_m$$

  - G is a geometric attenuation term (shadowing/masking).
- Or we can use explicitly defined microgeometry

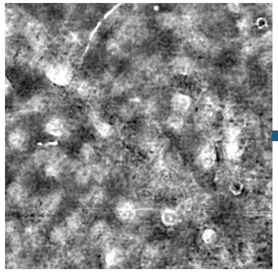

profilometry

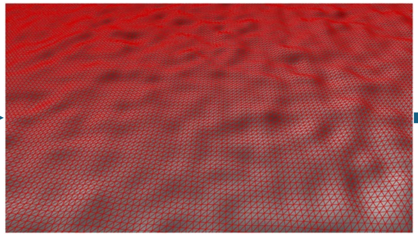

triangle mesh

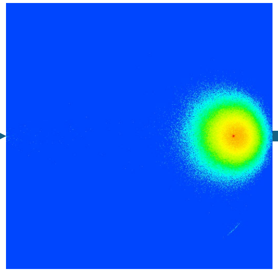

simulation

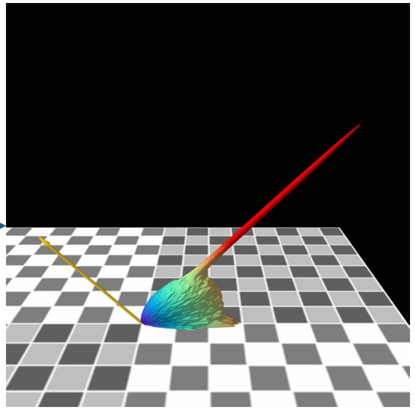

BSDF

Separability of optical effects

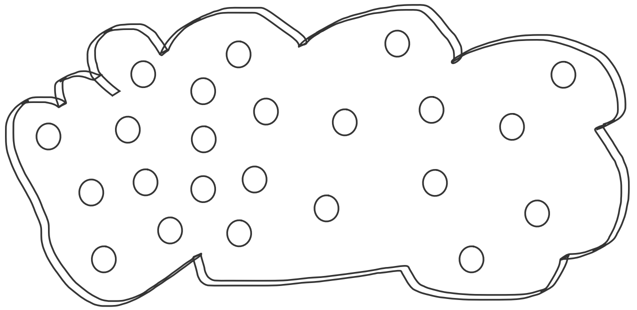

- Surface and volume [Ferrero *et al*. 2021. Optics Express]
  - Surface reflection is local $\delta(\boldsymbol{x}_r - \boldsymbol{x}_i)$ and shape (X) independent.
  - Volume effects are given by absorption and subsurface scattering.
- Subsurface scattering and absorption [Frisvad *et al*. 2007; 2012]
  - Scattering events are local and shape (X) independent.
  - Absorption and scattering lead to the probability that light follows a particular path in X.

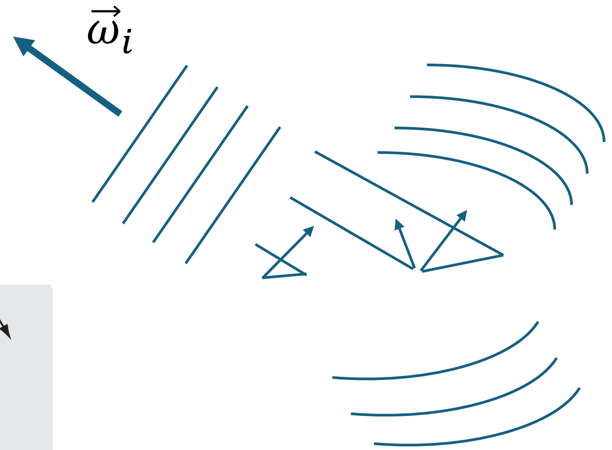

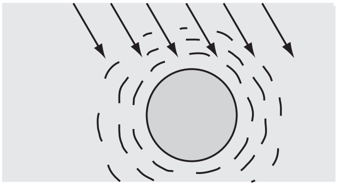

- Waves and rays [Falster *et al*. 2020. PG]
  - Wave effects are for coherent light in local geomety around the size of the wavelength.
  - Rays are sufficient for dealing with macroscopic paths in X.
- Coherence area and the Rayleigh criterion of optical smoothness
  - Coherence area limits the areal extent in which we would need to consider wave effects.
  - The Rayleigh criterion limits the resolution of the microgeometry that we would need for computing local bidirectional $(\vec{\omega}_i, \vec{\omega}_r)$ scattering/reflectance distributions.

Photo-render alignment

- For an object of known geometry on a planar surface, we can align a digital scene to a photo using silhouette matching if we have
  - Camera intrinsics (focal length / camera constant / field of view).
  - Simple lighting: point-like light source or diffuse lighting.
  - Segmentation of object, shadow, and background in the photo.
  - Approximate rotation of the object relative to the ground.
- Project silhouette edges onto the image plane.
- Use Blinn's projection shadows to find the light source position.

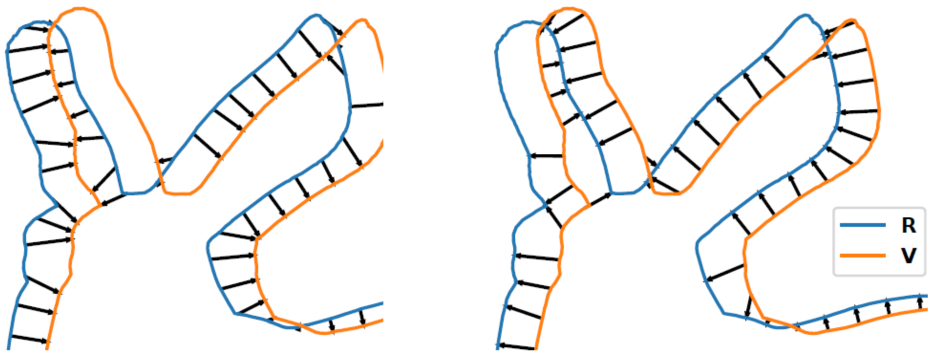

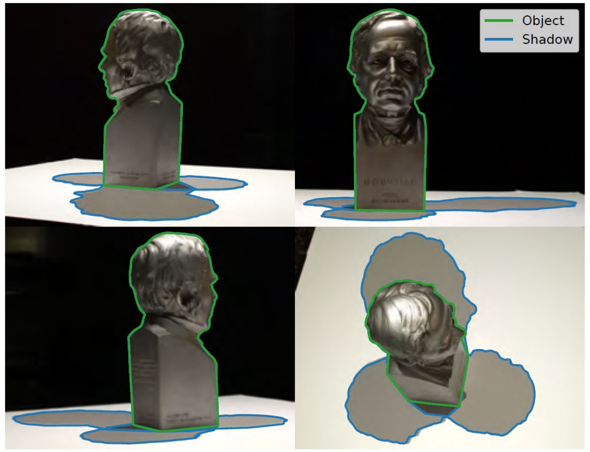

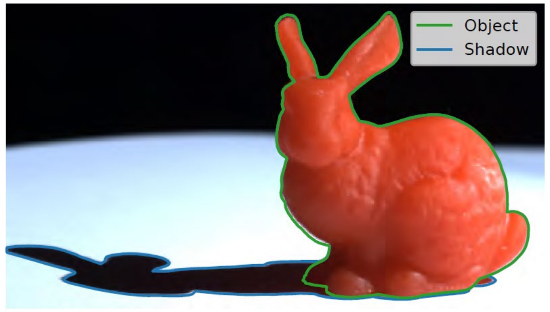

[Hannemose *et al*. 2020. Applied Optics]

Advancing macroscopic models (BSSRDF)

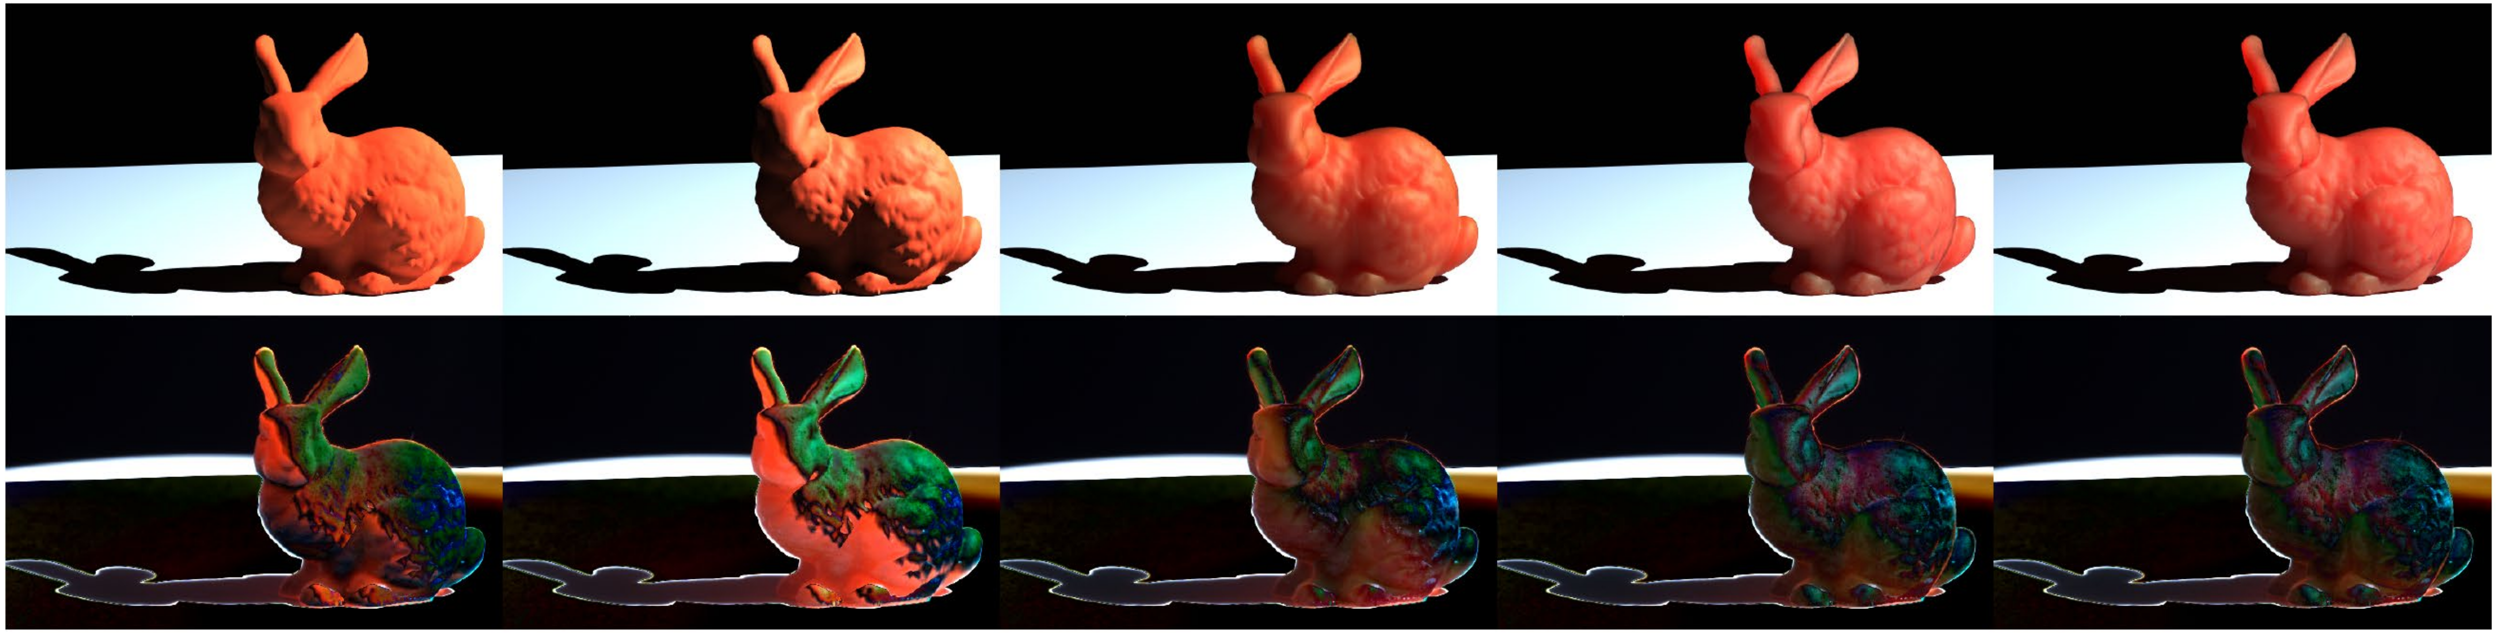

RMSE: 0.1152
SSIM: 0.8108
Lambertian

RMSE: 0.1237
SSIM: 0.7931
interfaced

RMSE: 0.1136
SSIM: 0.8145
standard SSS

RMSE: 0.1127
SSIM: 0.8177
directional SSS

RMSE: 0.1124
SSIM: 0.8180
SV roughness

[Hannemose *et al.* 2020]

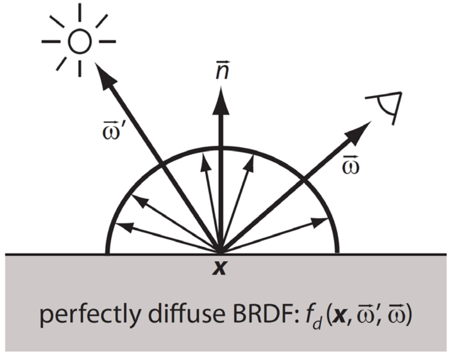

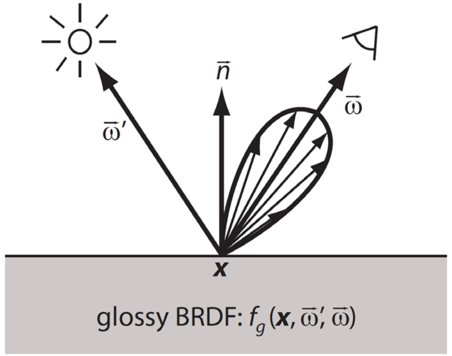

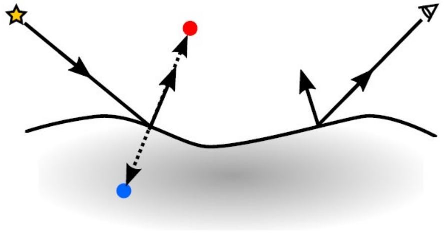

dipole model
[Jensen *et al.* 2001]

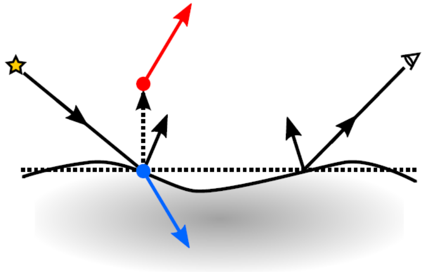

directional
dipole model
[Frisvad *et al.* 2014]

Importance of surface microstructure

3D printed translucent Stanford bunny

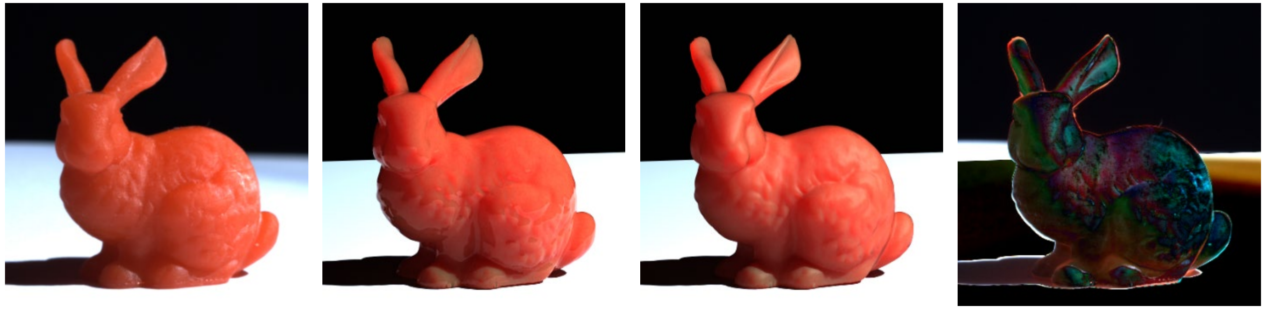

photograph | smooth | rough | abs diff × 2

Aluminium bust of H.C. Ørsted (3D scanned)

photo | smooth | rough | variation | abs diff × 2

3D scanned cupped angel 3D printed using transparent resin

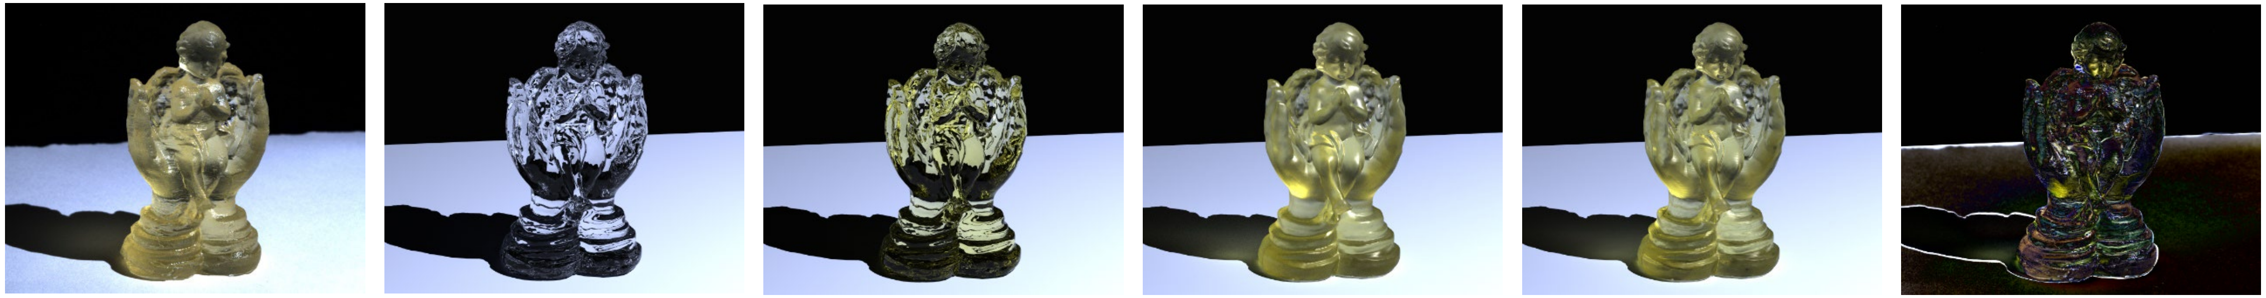

photograph of 3D print | rendering of 3D scan | absorption | random roughness | layered variation of roughness | absolute difference × 2

3D scanned figurine

[Hannemose *et al*. 2020. Applied Optics]

Estimating optical properties

- With photo-render alignment, we can use differentiable rendering to estimate optical properties.
- For heterogeneous volumes, this becomes highly challenging. The separations are difficult: surface or volume, absorption or scattering, wave effect or ray effect.
- Techniques exist for estimating optical properties using thin slabs:

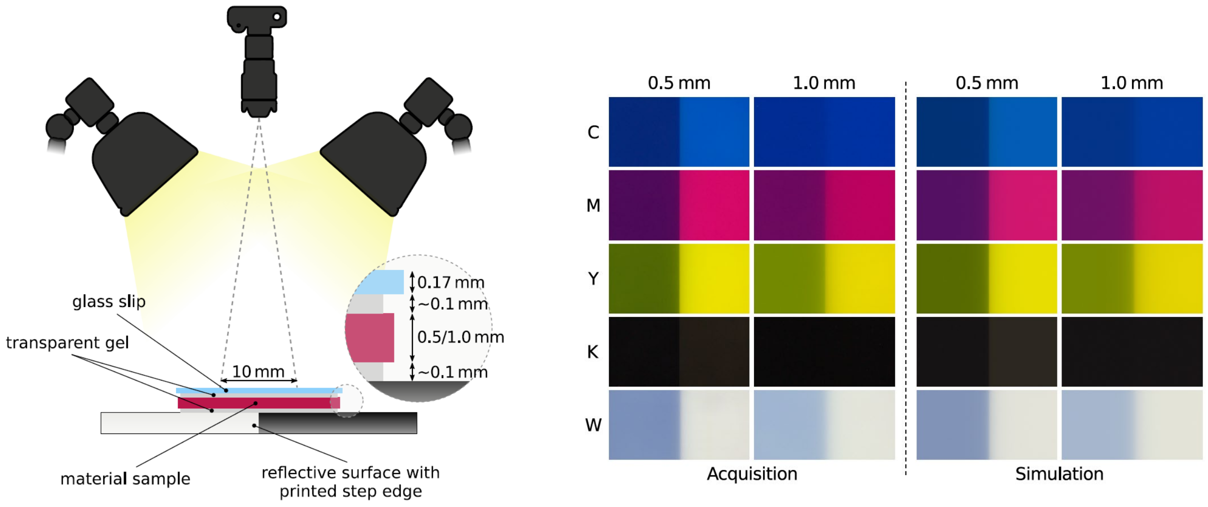

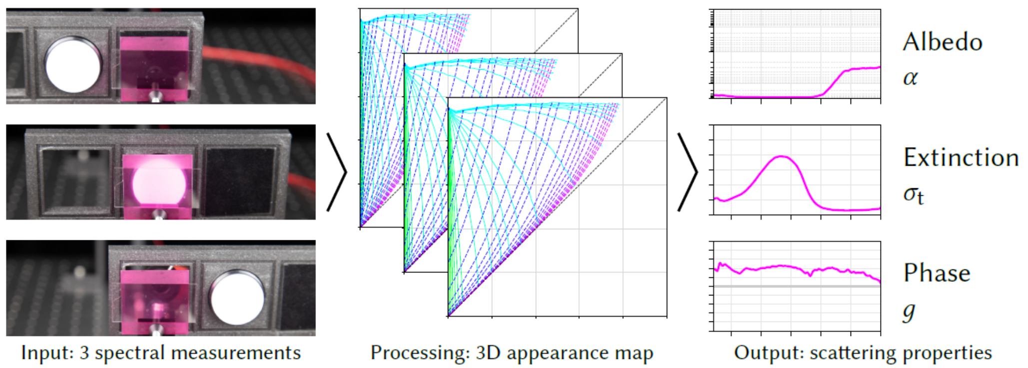

[Elek *et al.* 2021. Optics Express]

[Iser *et al.* 2022. SIGGRAPH Asia]

Alina Pranovich

- **Research project**
  ApPEARS: Appearance Printing European Advanced Research School
- **PhD project**
  Modeling appearance printing
- **Paper** (included in the following)

Digitizing the appearance of 3D printing materials using a spectrophotometer
Alina Pranovich, Morten Rieger Hannemose, Janus Nørtoft Jensen, Duc Minh Tran, Henrik Aanæs, Sasan Gooran, Daniel Nyström, Jeppe Revall Frisvad
Sensors 24(21), Article 7025. October 2024.

ApPEARS
APPEARANCE PRINTING
European Advanced Research School

3D printing is an excellent tool for validation

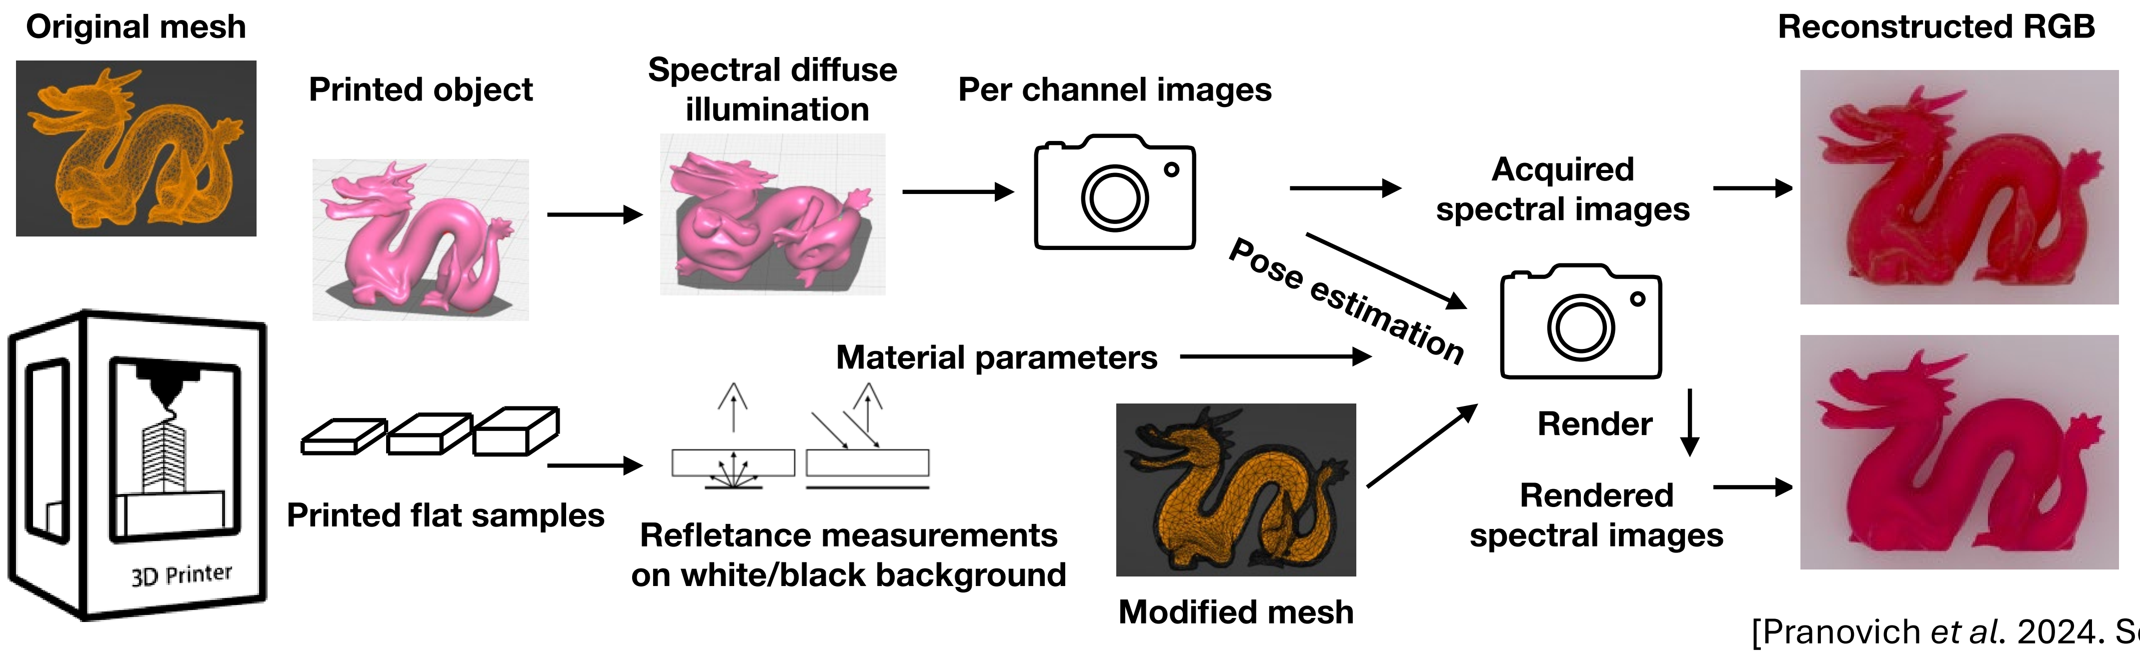

[Pranovich *et al*. 2024. Sensors]

- Suppose we have a way of estimating optical properties from thin slabs.
- We can use photo-render alignment to validate the correctness of the estimated optical properties.

Estimation and validation of optical properties

- Based on scattering in isotropic plane-parallel media with a rough surface, we built an **analytic model** representing a **spectrophotometer**.
- This enables us to estimate **spectral optical properties** based on a collection of **thin slab samples**.
- We then **3D print an object** with non-trivial geometry (the Stanford dragon) and use **spectral photo-render comparison under diffuse lighting** to test the correctness of our estimated optical properties.

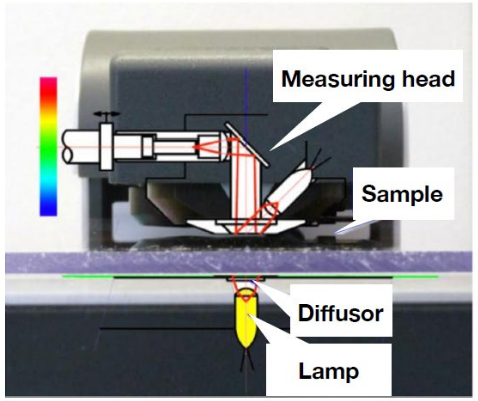

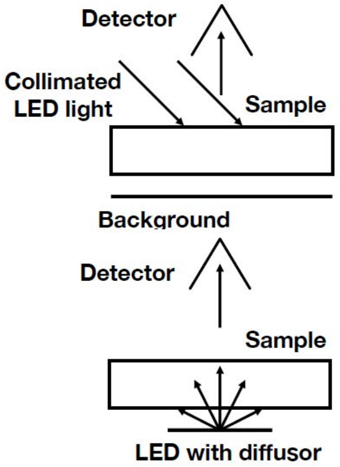

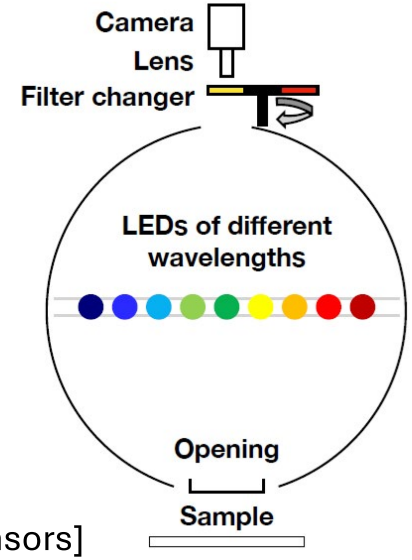

[Pranovich *et al.* 2024. Sensors]

Estimation and validation of optical properties

- Estimation (dashed curves based on appearance maps [Iser *et al.* 2022])

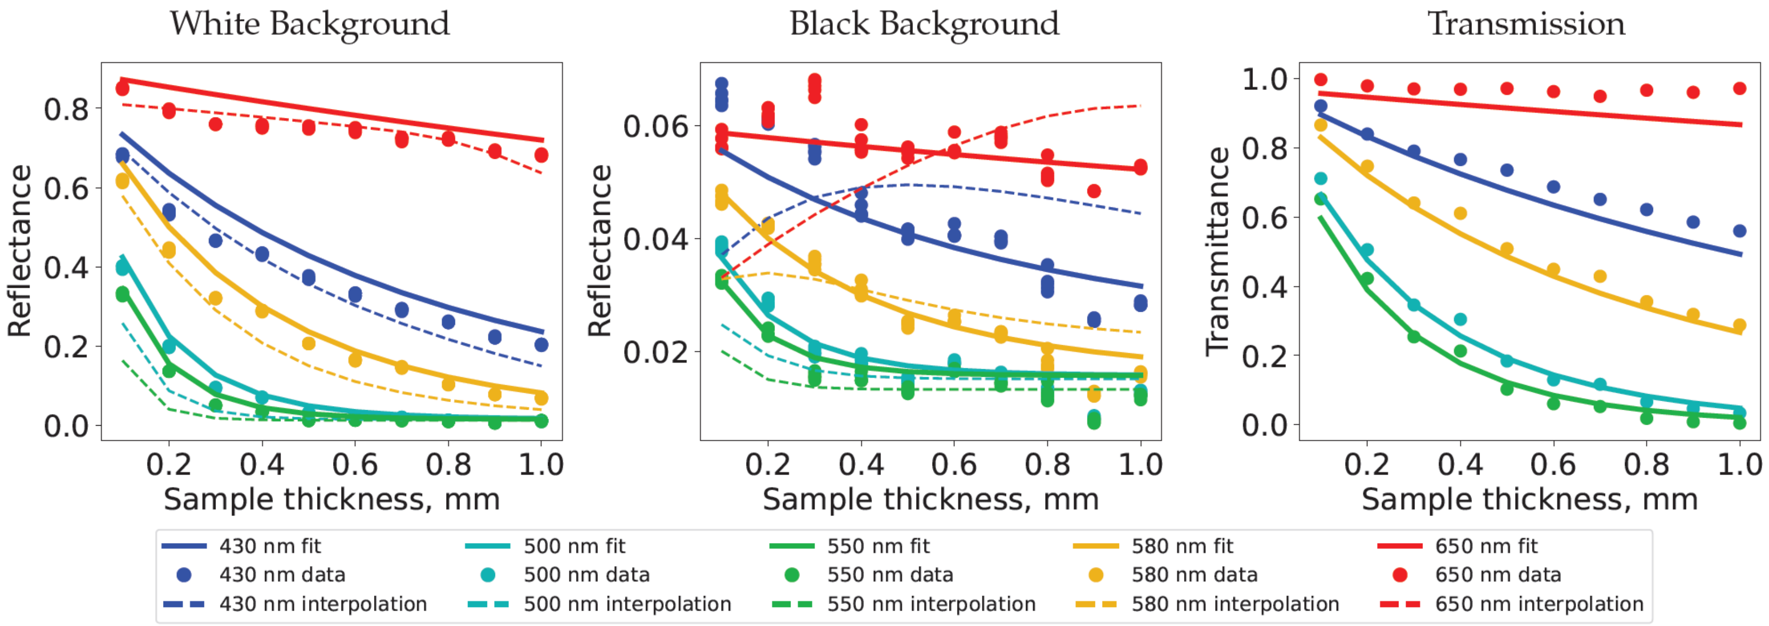

- Validation (assessment of correctness)

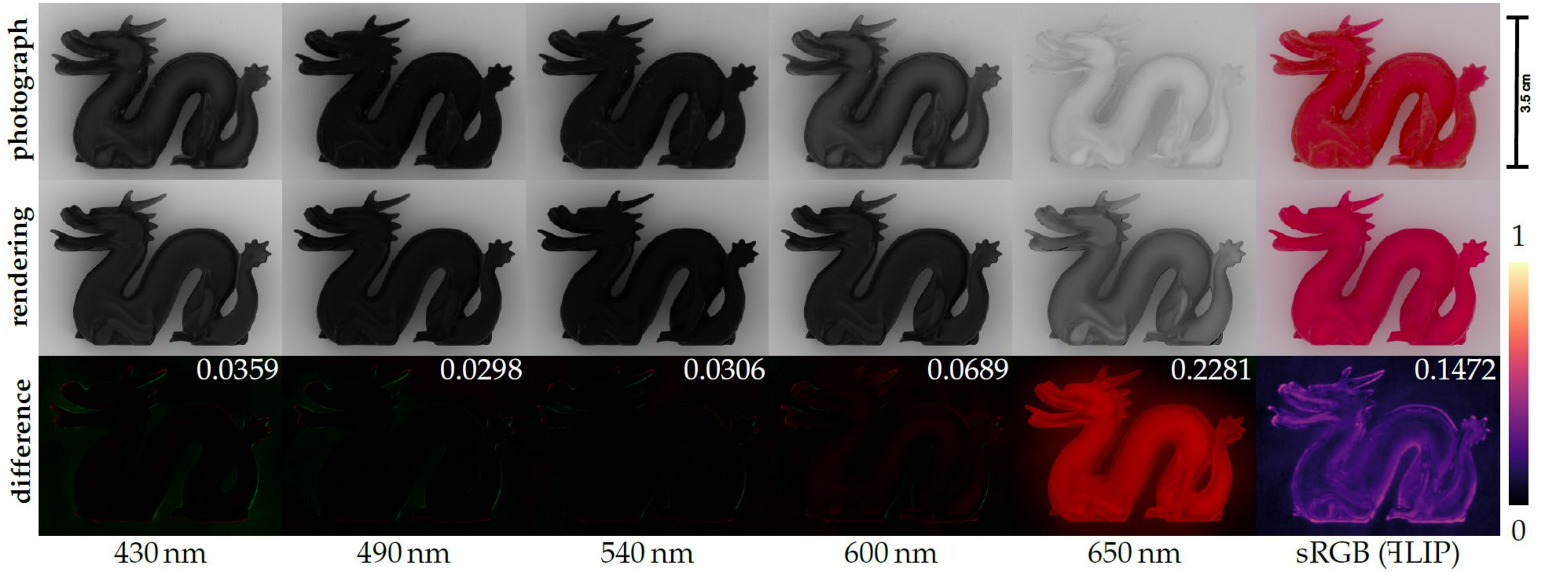

[Pranovich *et al.* 2024. Sensors]

Photo-render comparisons (sRGB reconstructions)

measure

model

1 mm slabs of 3D printer primary inks
on a white background

- Testing how closely we can match the color of thin samples on a white background (see above) seems insufficient in terms of testing the predictive rendering capabilities of the estimated optical properties.

[Pranovich *et al.* 2024. Sensors]

Changing the lighting conditions

- Photo of objects made with different mixes of Vero Red and Vero White (and with white infill).

- Renderings of their digital twins in different photographed environments.

Duc Minh Tran

- **Research project**
  BxDiff: new quantities for the measurement of appearance
- **PhD project**
  Rendering of objects with measured translucent appearance
- **Paper** (included in the following)

Digitizing translucent object appearance by validating computed optical properties
Duc Minh Tran, Mark Bo Jensen, Pablo Santafé-Gabarda, Stefan Källberg, Alejandro Ferrero, Morten Rieger Hannemose, Jeppe Revall Frisvad
Applied Optics 63(16), pp. 4317-4331. June 2024.

The EMPIR initiative is co-funded by the European Union's Horizon 2020 research and innovation programme and the EMPIR Participating States

JRP 18SIB03

Simulation based on microgeometry

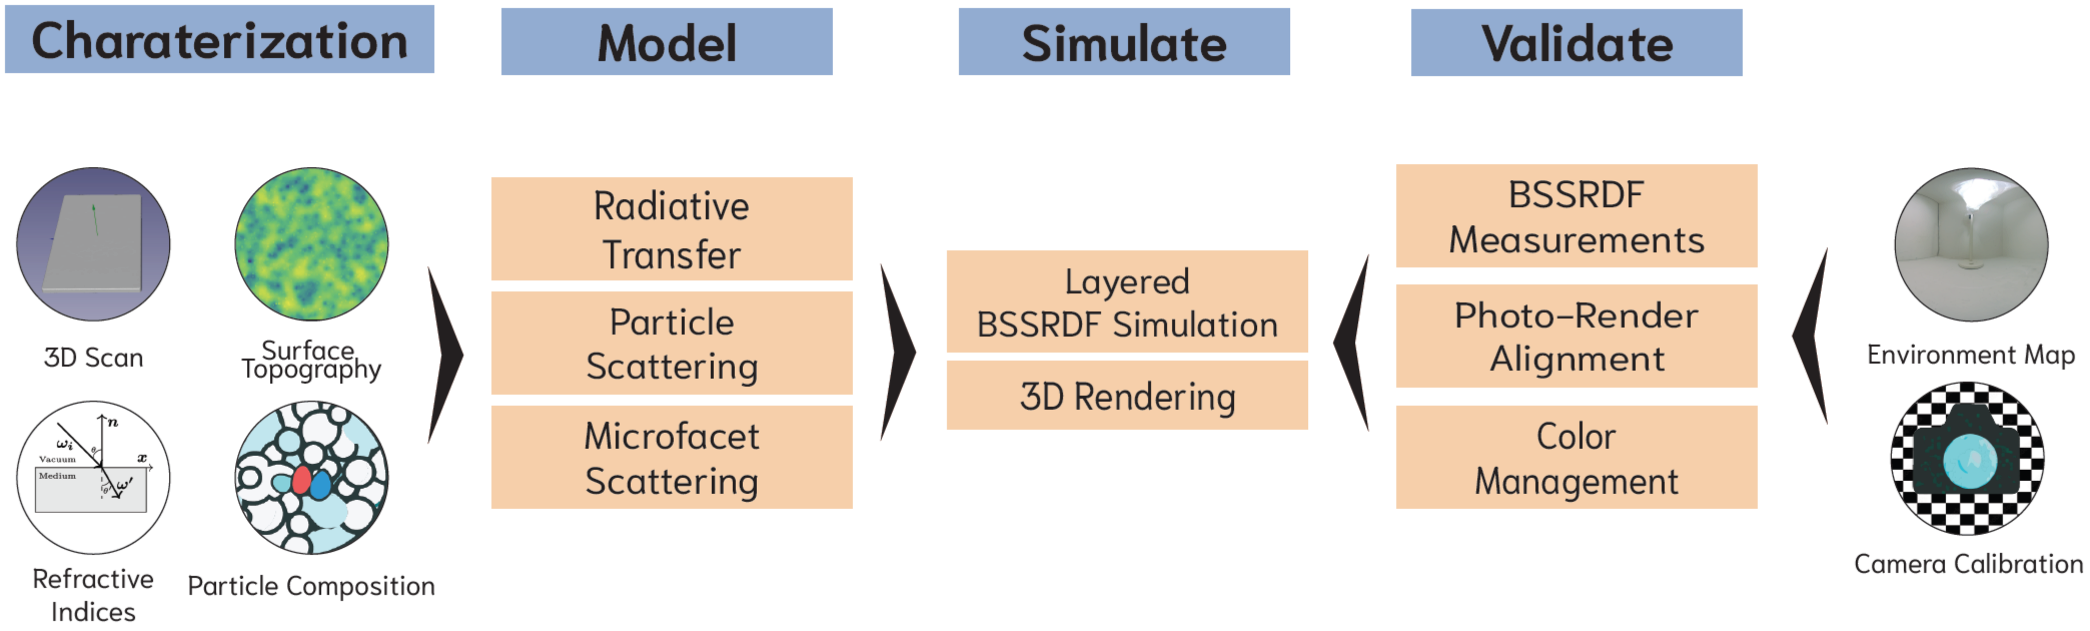

[Tran *et al*. 2024. Applied Optics]

- Based on data from BxDiff
  https://bxdiff.cmi.gov.cz/

EMPIR EURAMET
The EMPIR initiative is co-funded by the European Union's Horizon 2020 research and innovation programme and the EMPIR Participating States

BxDiff
JRP 18SIB03

Sample characterization

1. Host medium refractive index.
2. Particle type, volume fraction (or wt.-% or density), refractive index.
3. Particle size distribution (at least mean particle size).
4. Sample surface geometry (3D scan or CAD file).
5. Surface topography (profilometry scan).

- Info on subsurface particles can be obtained from a micro-CT scan.
- Properties of individual particles can be obtained from interferometry.

Rendering of Translucent Materials

- Volume rendering: solve the radiative transfer equation (RTE).

$$(\vec{\omega}\cdot\nabla)L(\boldsymbol{x},\vec{\omega}) = \overbrace{-\sigma_t L(\boldsymbol{x},\vec{\omega})}^{\text{absorption and out-scattering}} + \overbrace{\sigma_s \int_{4\pi} p(\vec{\omega}',\vec{\omega})L(\boldsymbol{x},\vec{\omega}')\,\mathrm{d}\omega'}^{\text{in-scattering}}$$

input: scattering properties

- General solution: path tracing (Monte Carlo integration).

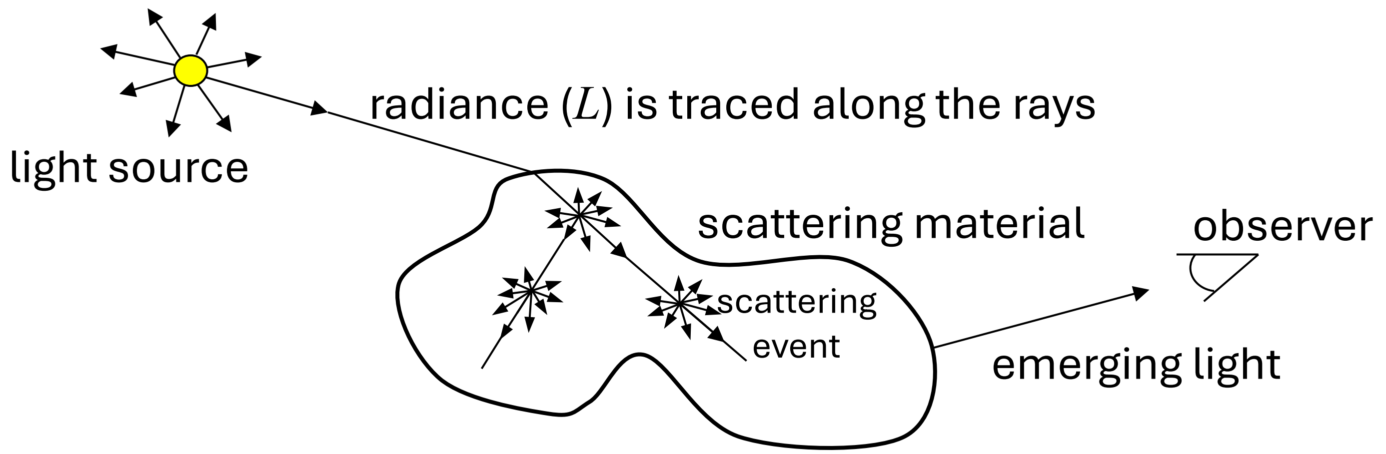

- Predicting appearance: requires a model for computing scattering properties (Lorenz-Mie theory is an option).

Appearance based on scattering by particles

- Lorenz-Mie theory describes the scattering of plane waves by smooth spherical particles (of arbitrary size).

Phase function of a 20 nm casein micelle

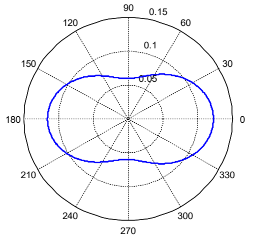

Phase function of a 1 μm fat globule

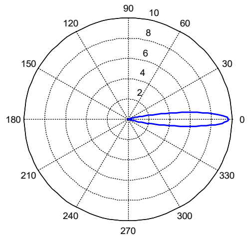

Phase function of a 1 μm fat globule - close-up

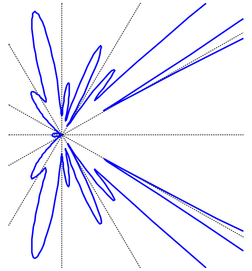

- Assume that particles scatter light independently (decoupling).
- Then integration over particle size distributions provide the scattering properties used in a rendering. Examples:

seawater

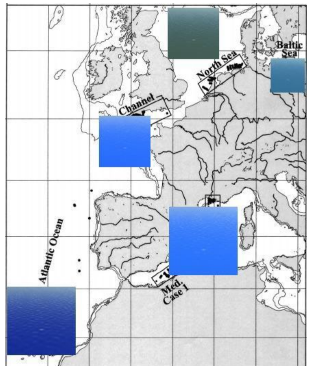

algae in sea ice

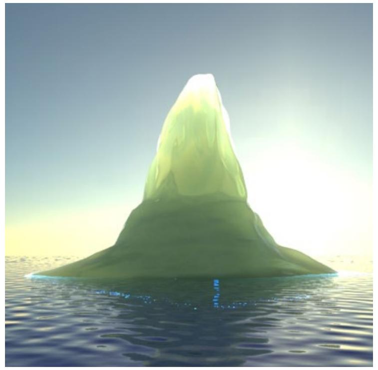

unfiltered apple juice

photo render

BSSRDF measurements from CSIC

- Goniometric measurements of subsurface scattering.
- Reasonable but not perfect match with measurements.
- We can get a very good match by adjusting our input parameters.
- Most likely the input parameters were imprecise.

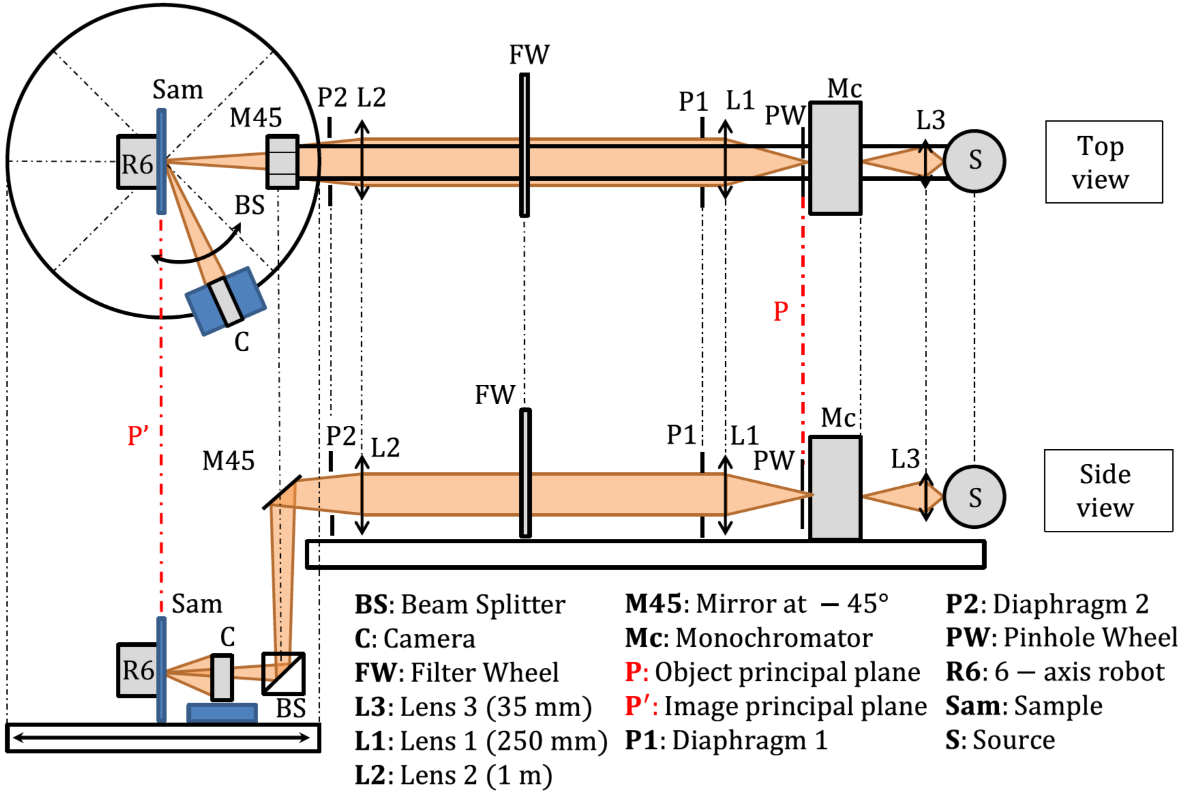

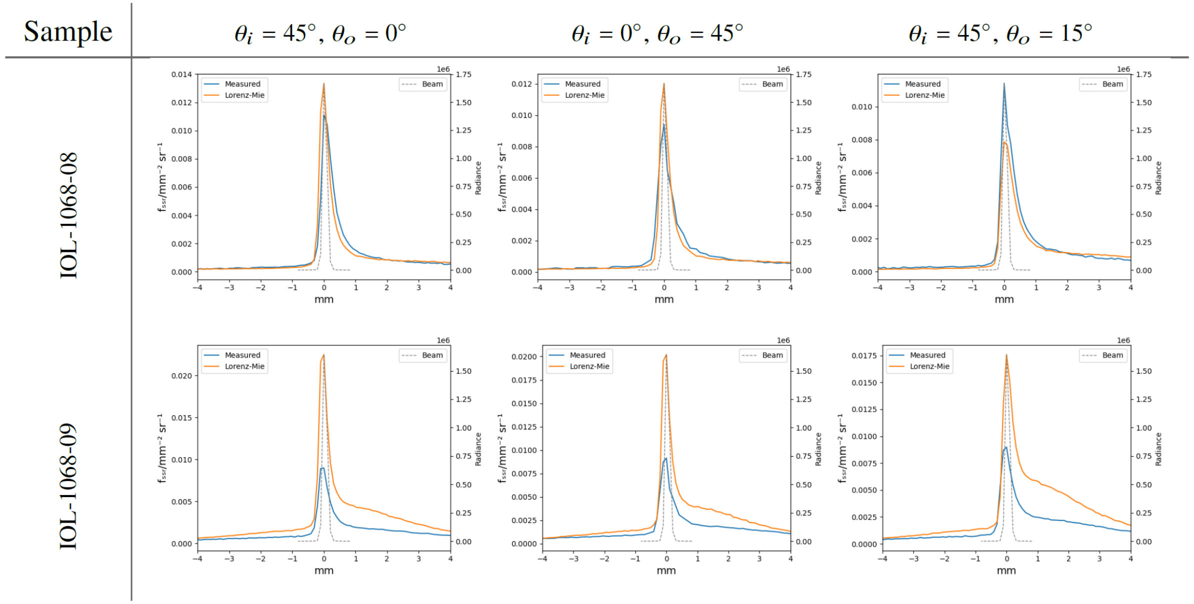

Simulation

- Controlled lighting.

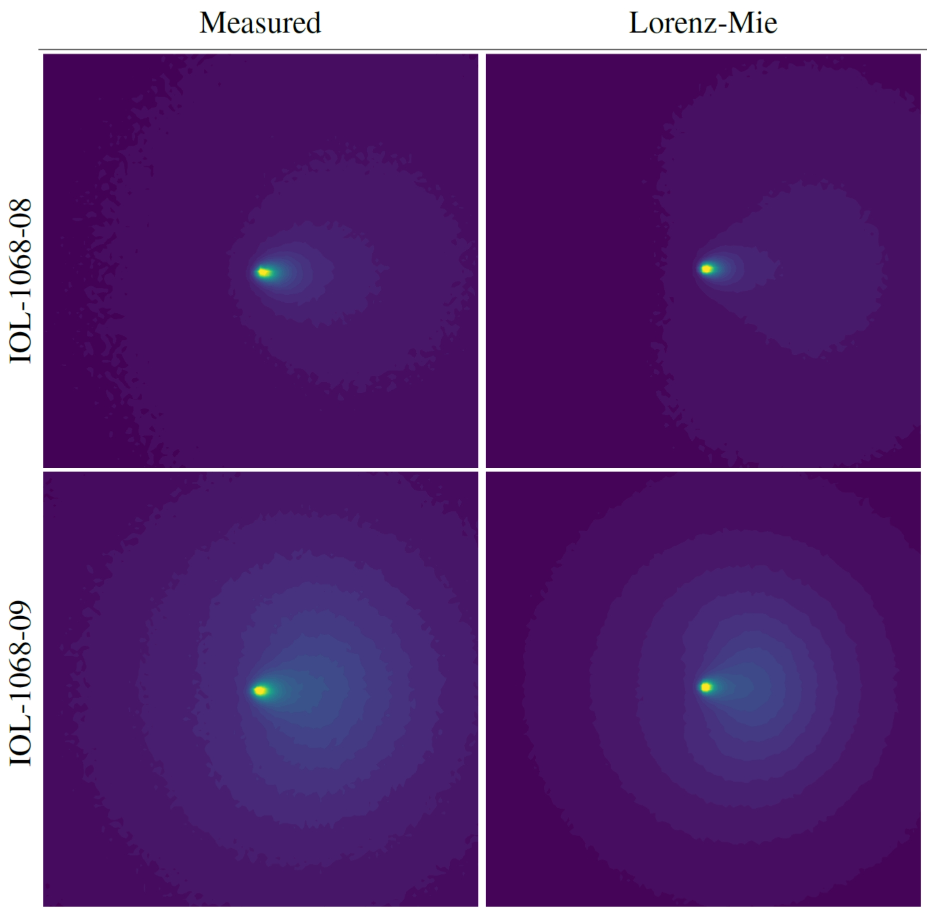

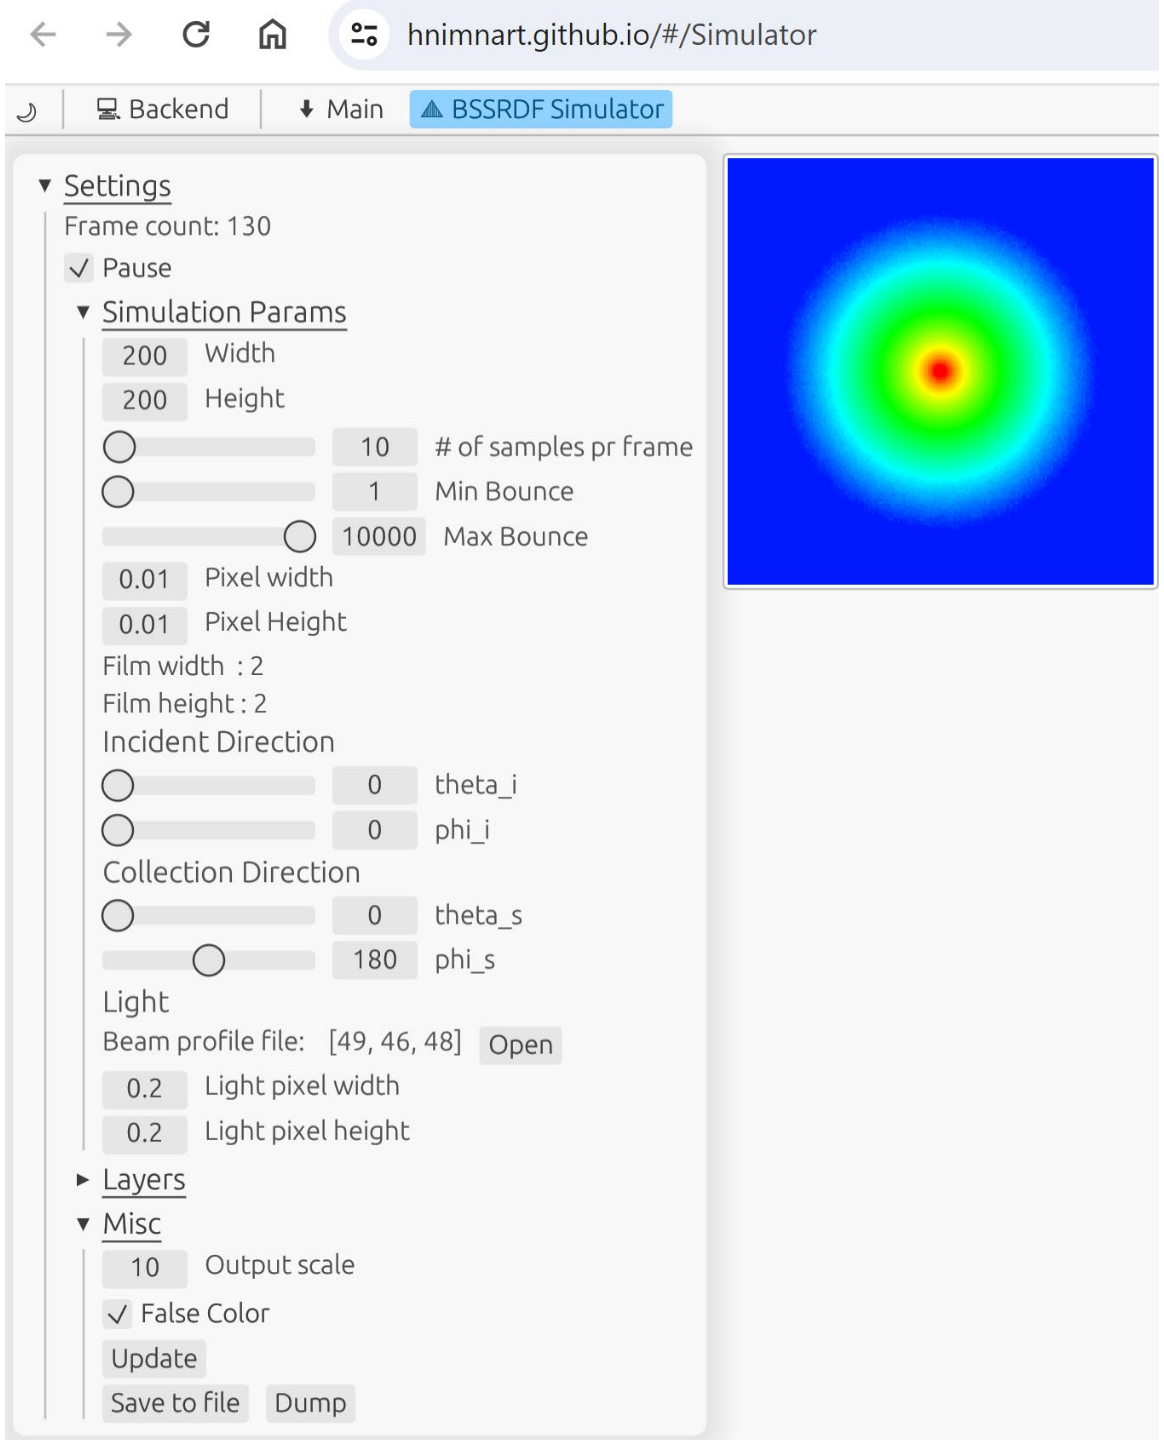

Adjustment of parameters

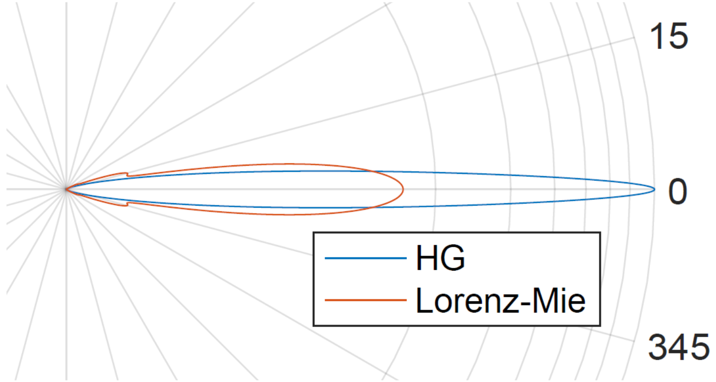

- We found an approximate phase function like Henyey-Greenstein is too inaccurate.

- Small adjustment of the particle size adjusts the phase function enough to improve the fit.

- Adjusting the volume fraction also important.

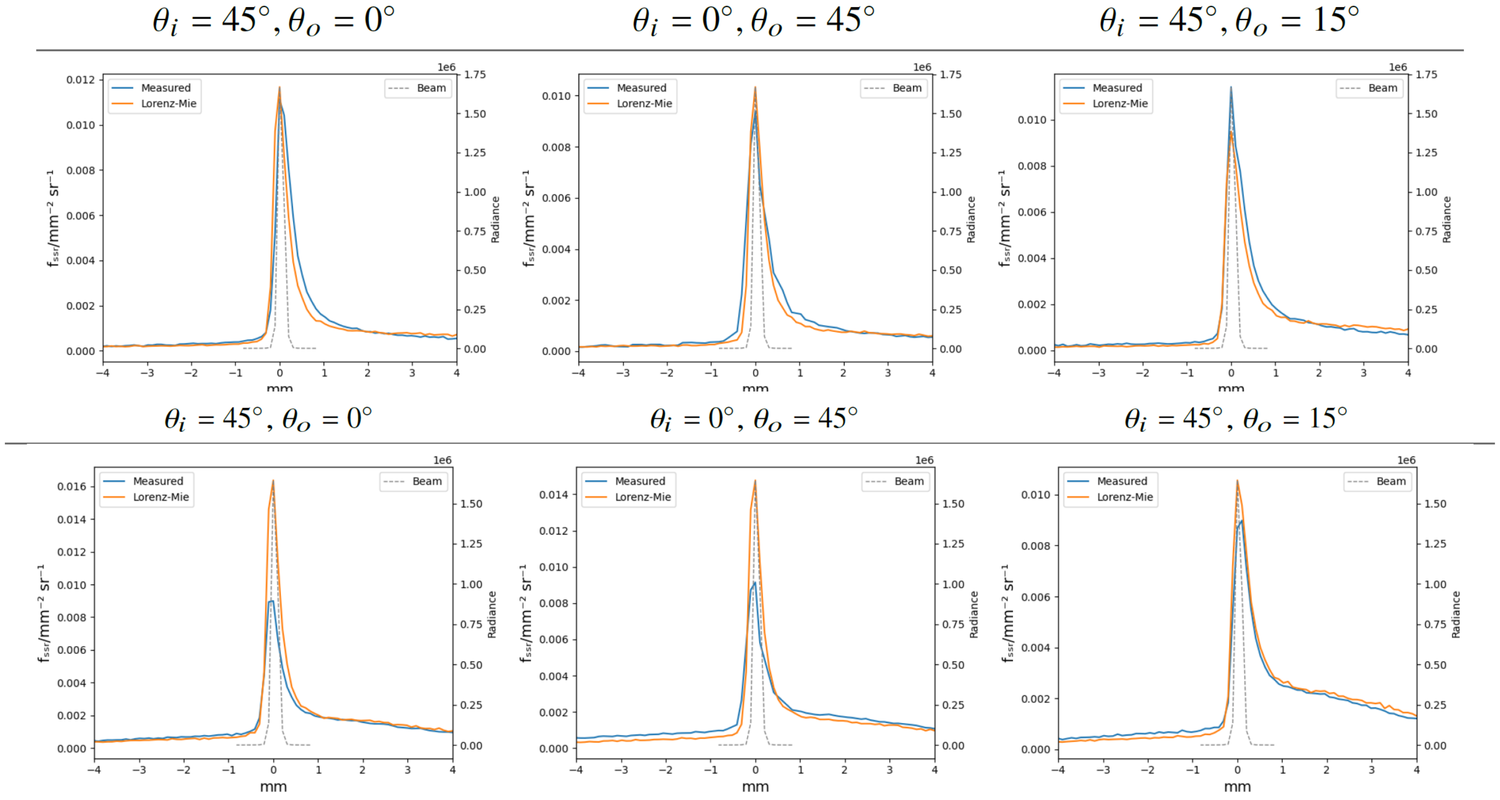

Photo-render comparison

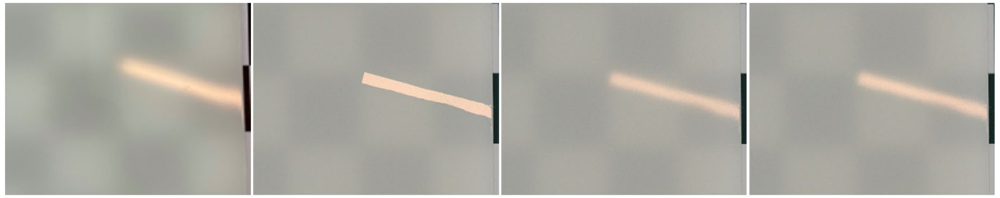

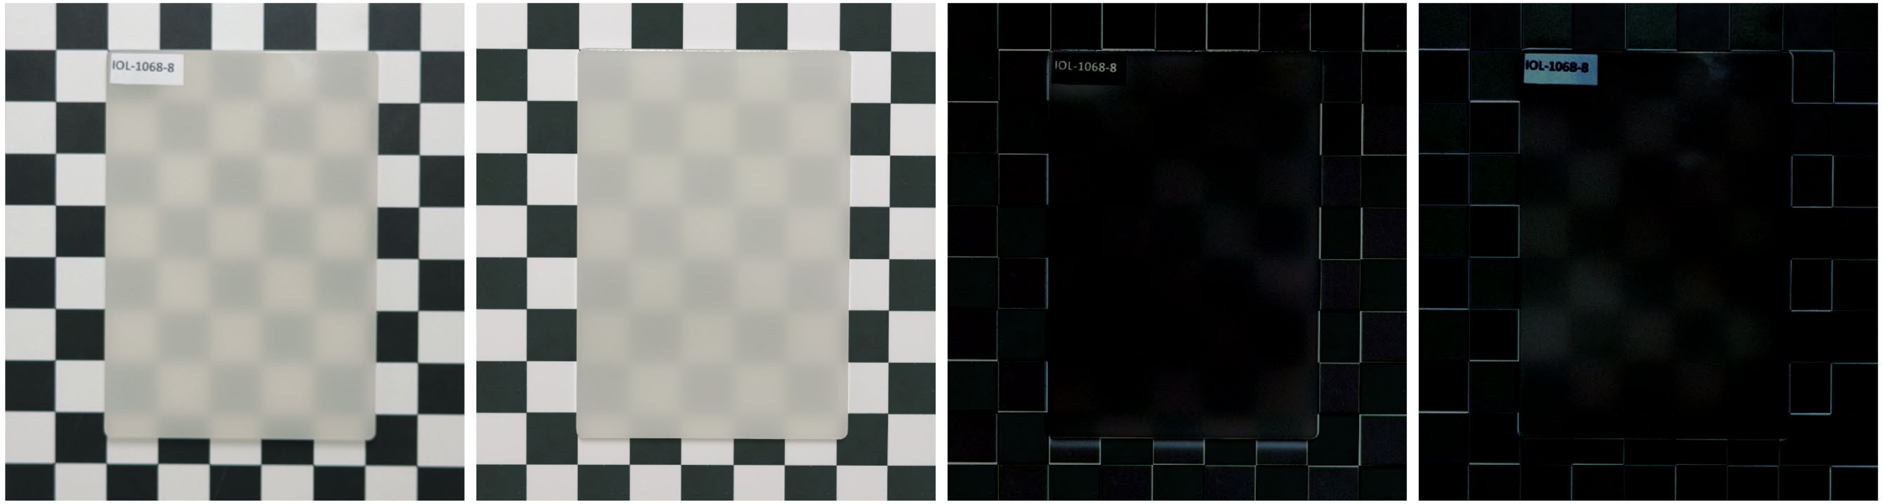

The importance of accounting for surface roughness.

IOL-1068-08

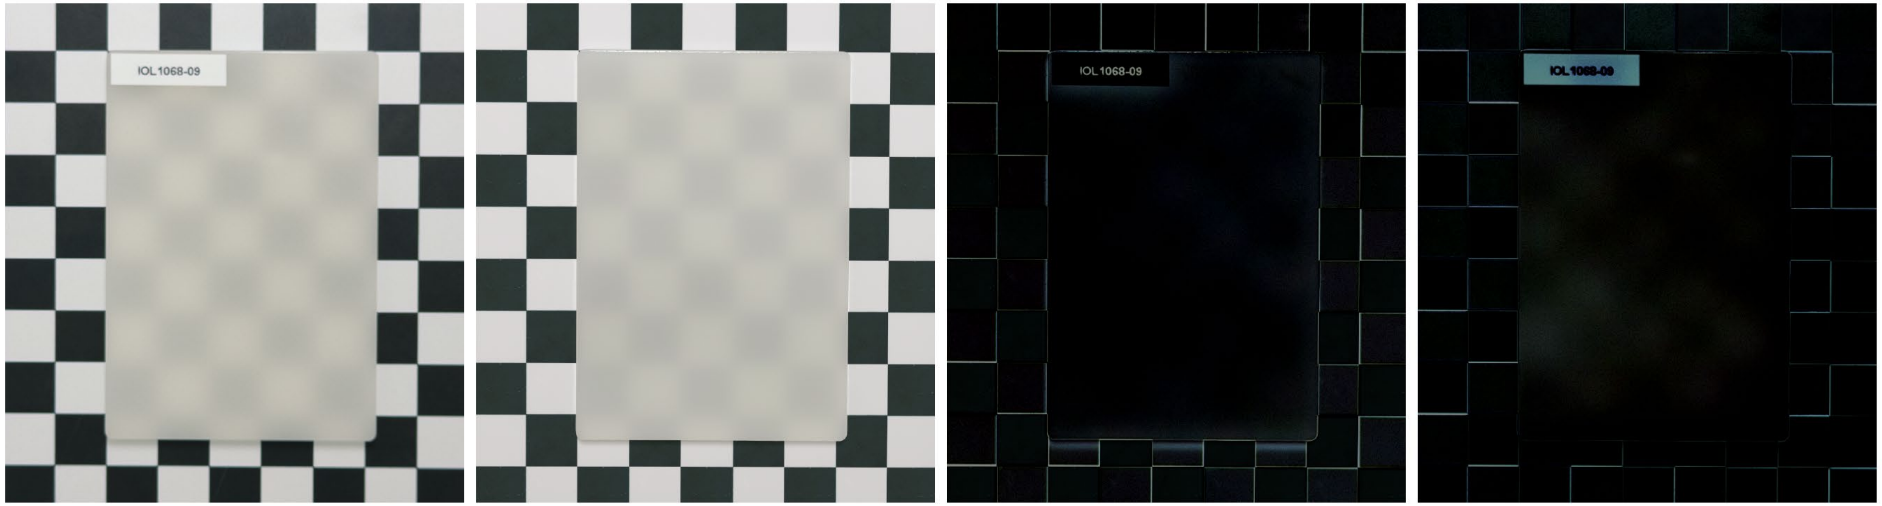

IOL-1068-09

Good match after adjustment of parameters

Predictive rendering

- We are now rendering different samples to specify production of a set of samples that spans the scale of perceived translucency.
- The research challenge is to create a metric for perceived translucency.

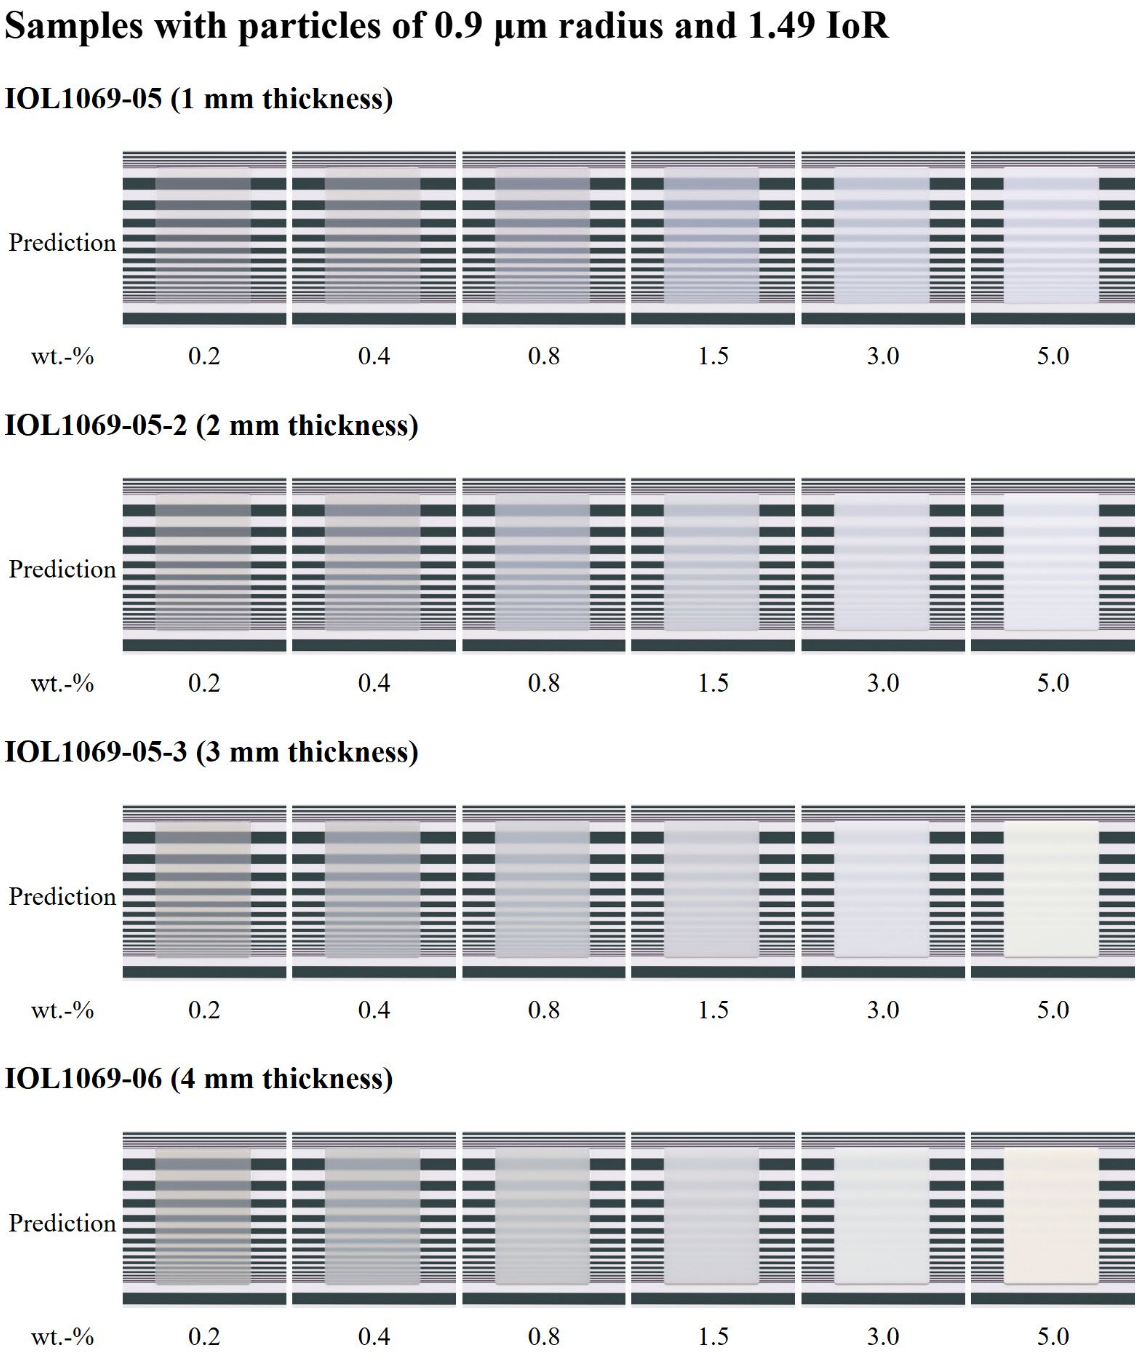

Thank you for your attention

[Luongo *et al.* 2020. CGF]