Introduction to Computer Graphics

GAMES101, Lingqi Yan, UC Santa Barbara

Lecture 18:
Advanced Topics in Rendering

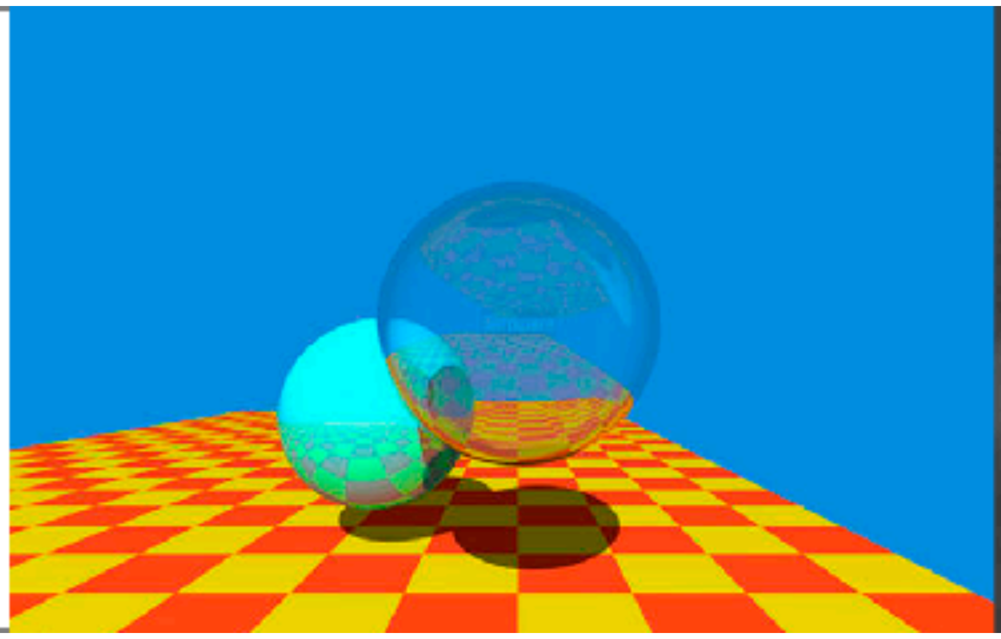

Announcements

- Homework 7 will be released soon
- Final project timeline
  - [Apr 14] Ideas will be released next Tuesday
  - [Apr 19] Submit your proposal one week later
  - [May 5] Submit your work
- Final project logistics
  - Work on Graphics topics, write code on your own
- Today's lecture
  - Advanced (?) light transport and materials
  - A lot, but extremely high-level. Mostly FYI.

Advanced Light Transport

Advanced Light Transport

- Unbiased light transport methods
  - Bidirectional path tracing (BDPT)
  - Metropolis light transport (MLT)
- Biased light transport methods
  - Photon mapping
  - Vertex connection and merging (VCM)
- Instant radiosity (VPL / many light methods)

Biased vs. Unbiased Monte Carlo Estimators

- An unbiased Monte Carlo technique does not have any systematic error
  - The expected value of an unbiased estimator will always be the correct value, no matter how many samples are used
- Otherwise, biased
  - One special case, the expected value converges to the correct value as infinite #samples are used — consistent
- We'll look again at this page after introducing Photon Mapping

Bidirectional Path Tracing (BDPT)

- Recall: a path connects the camera and the light
- BDPT
  - Traces sub-paths from both the camera and the light
  - Connects the end points from both sub-paths

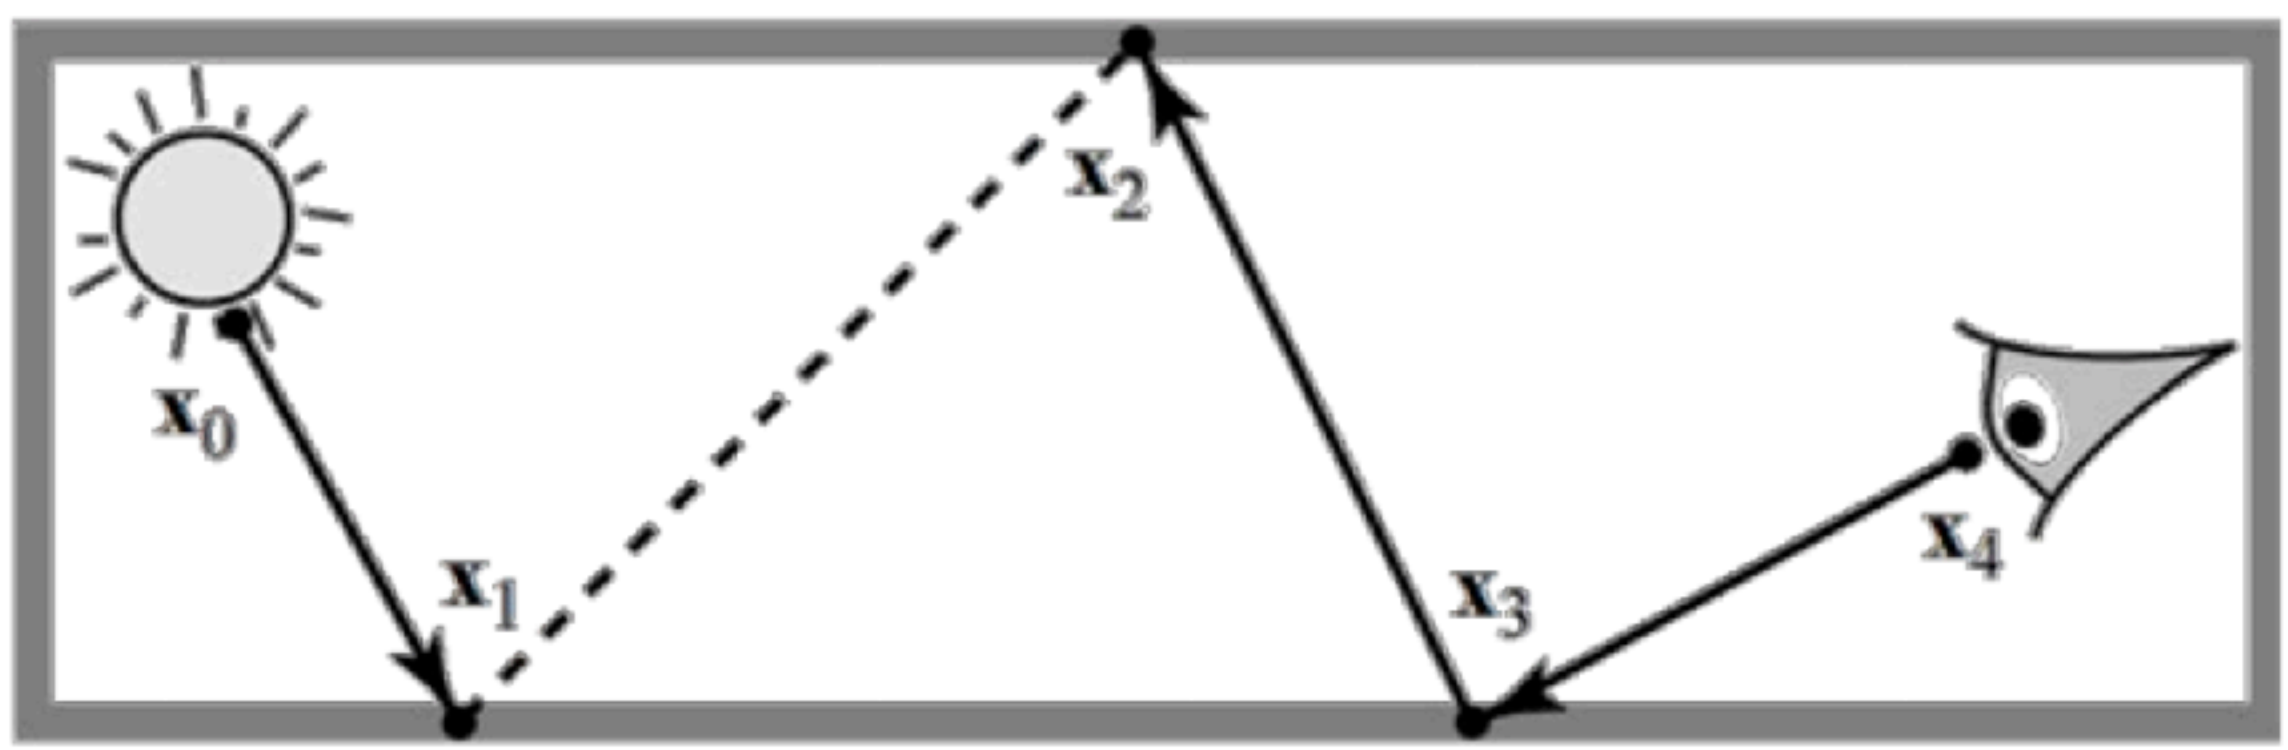

[Veach 1997]

Bidirectional Path Tracing (BDPT)

- Suitable if the light transport is complex on the light's side
- Difficult to implement & quite slow

(a) Path tracer, 32 samples/pixel (b) Bidirectional path tracer, 32 samples/pixel

Metropolis Light Transport (MLT)

- A Markov Chain Monte Carlo (MCMC) application
  - Jumping from the current sample to the next with some PDF
- Very good at locally exploring difficult light paths
- Key idea
  - Locally perturb an existing path to get a new path

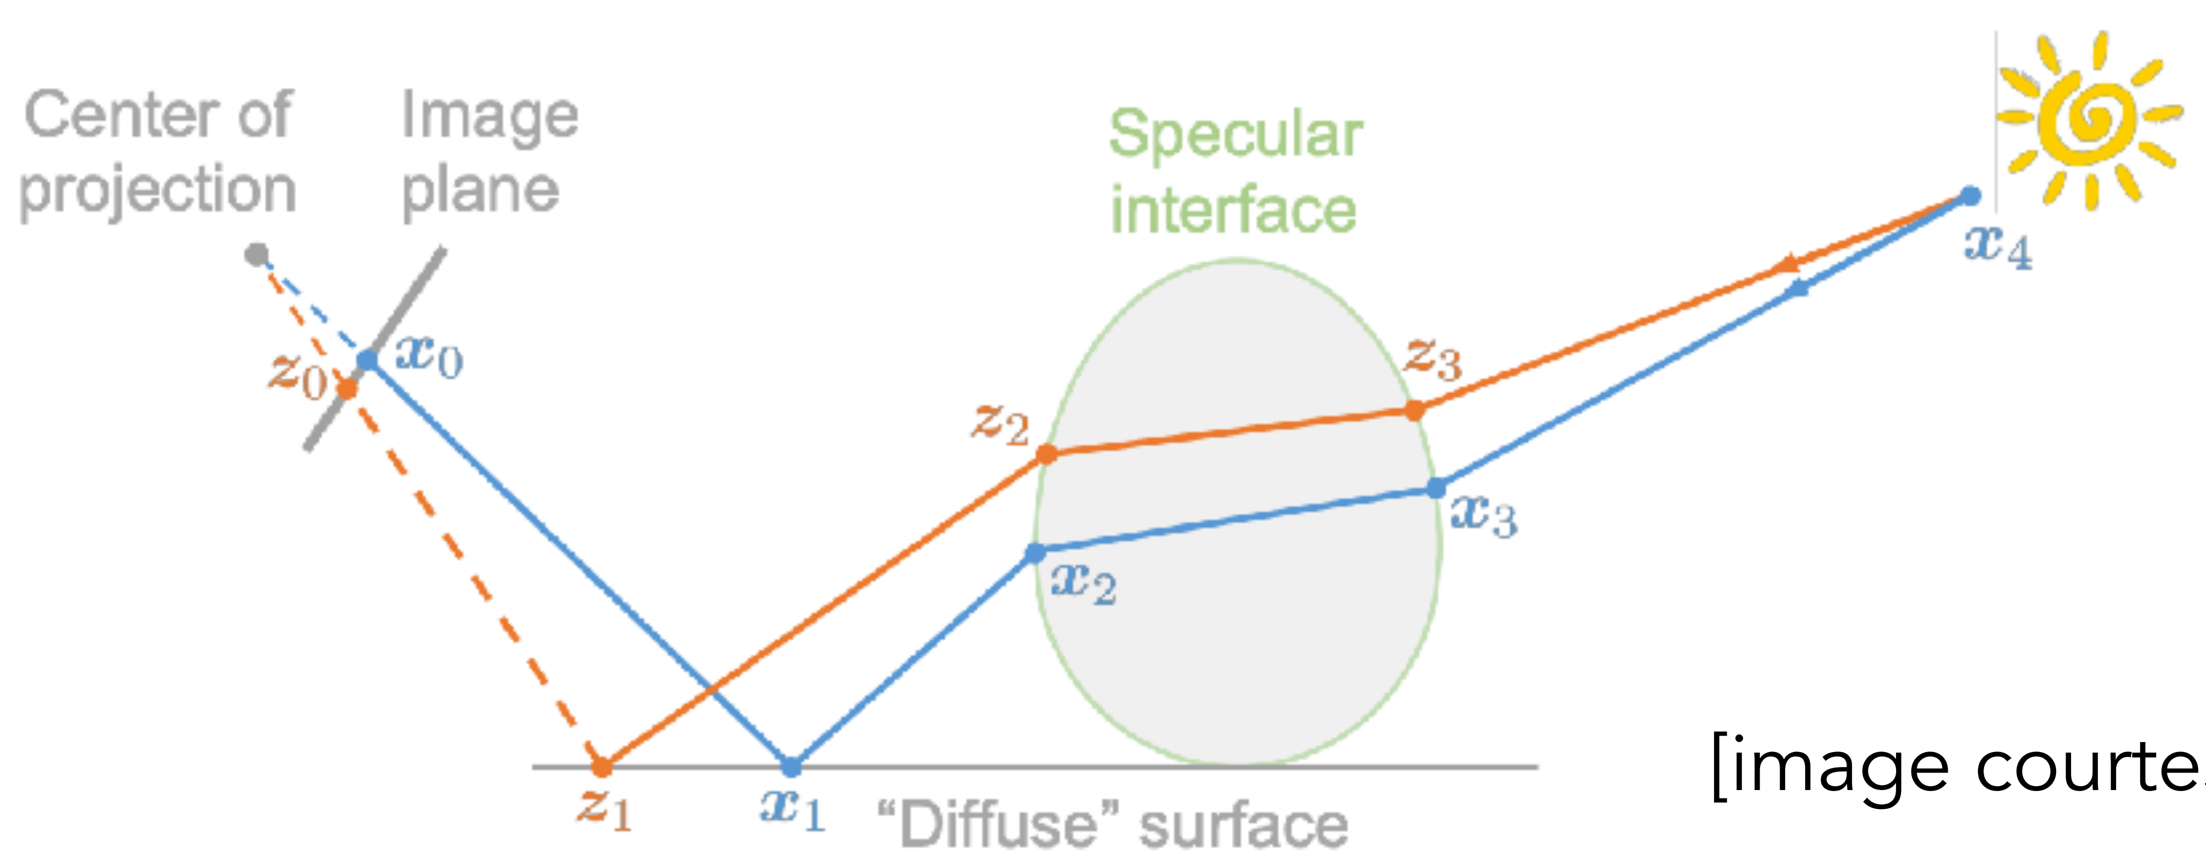

[image courtesy of Shuang Zhao]

Metropolis Light Transport (MLT) — Pros

- Works great with difficult light paths
- Also unbiased

BDPT

MLT

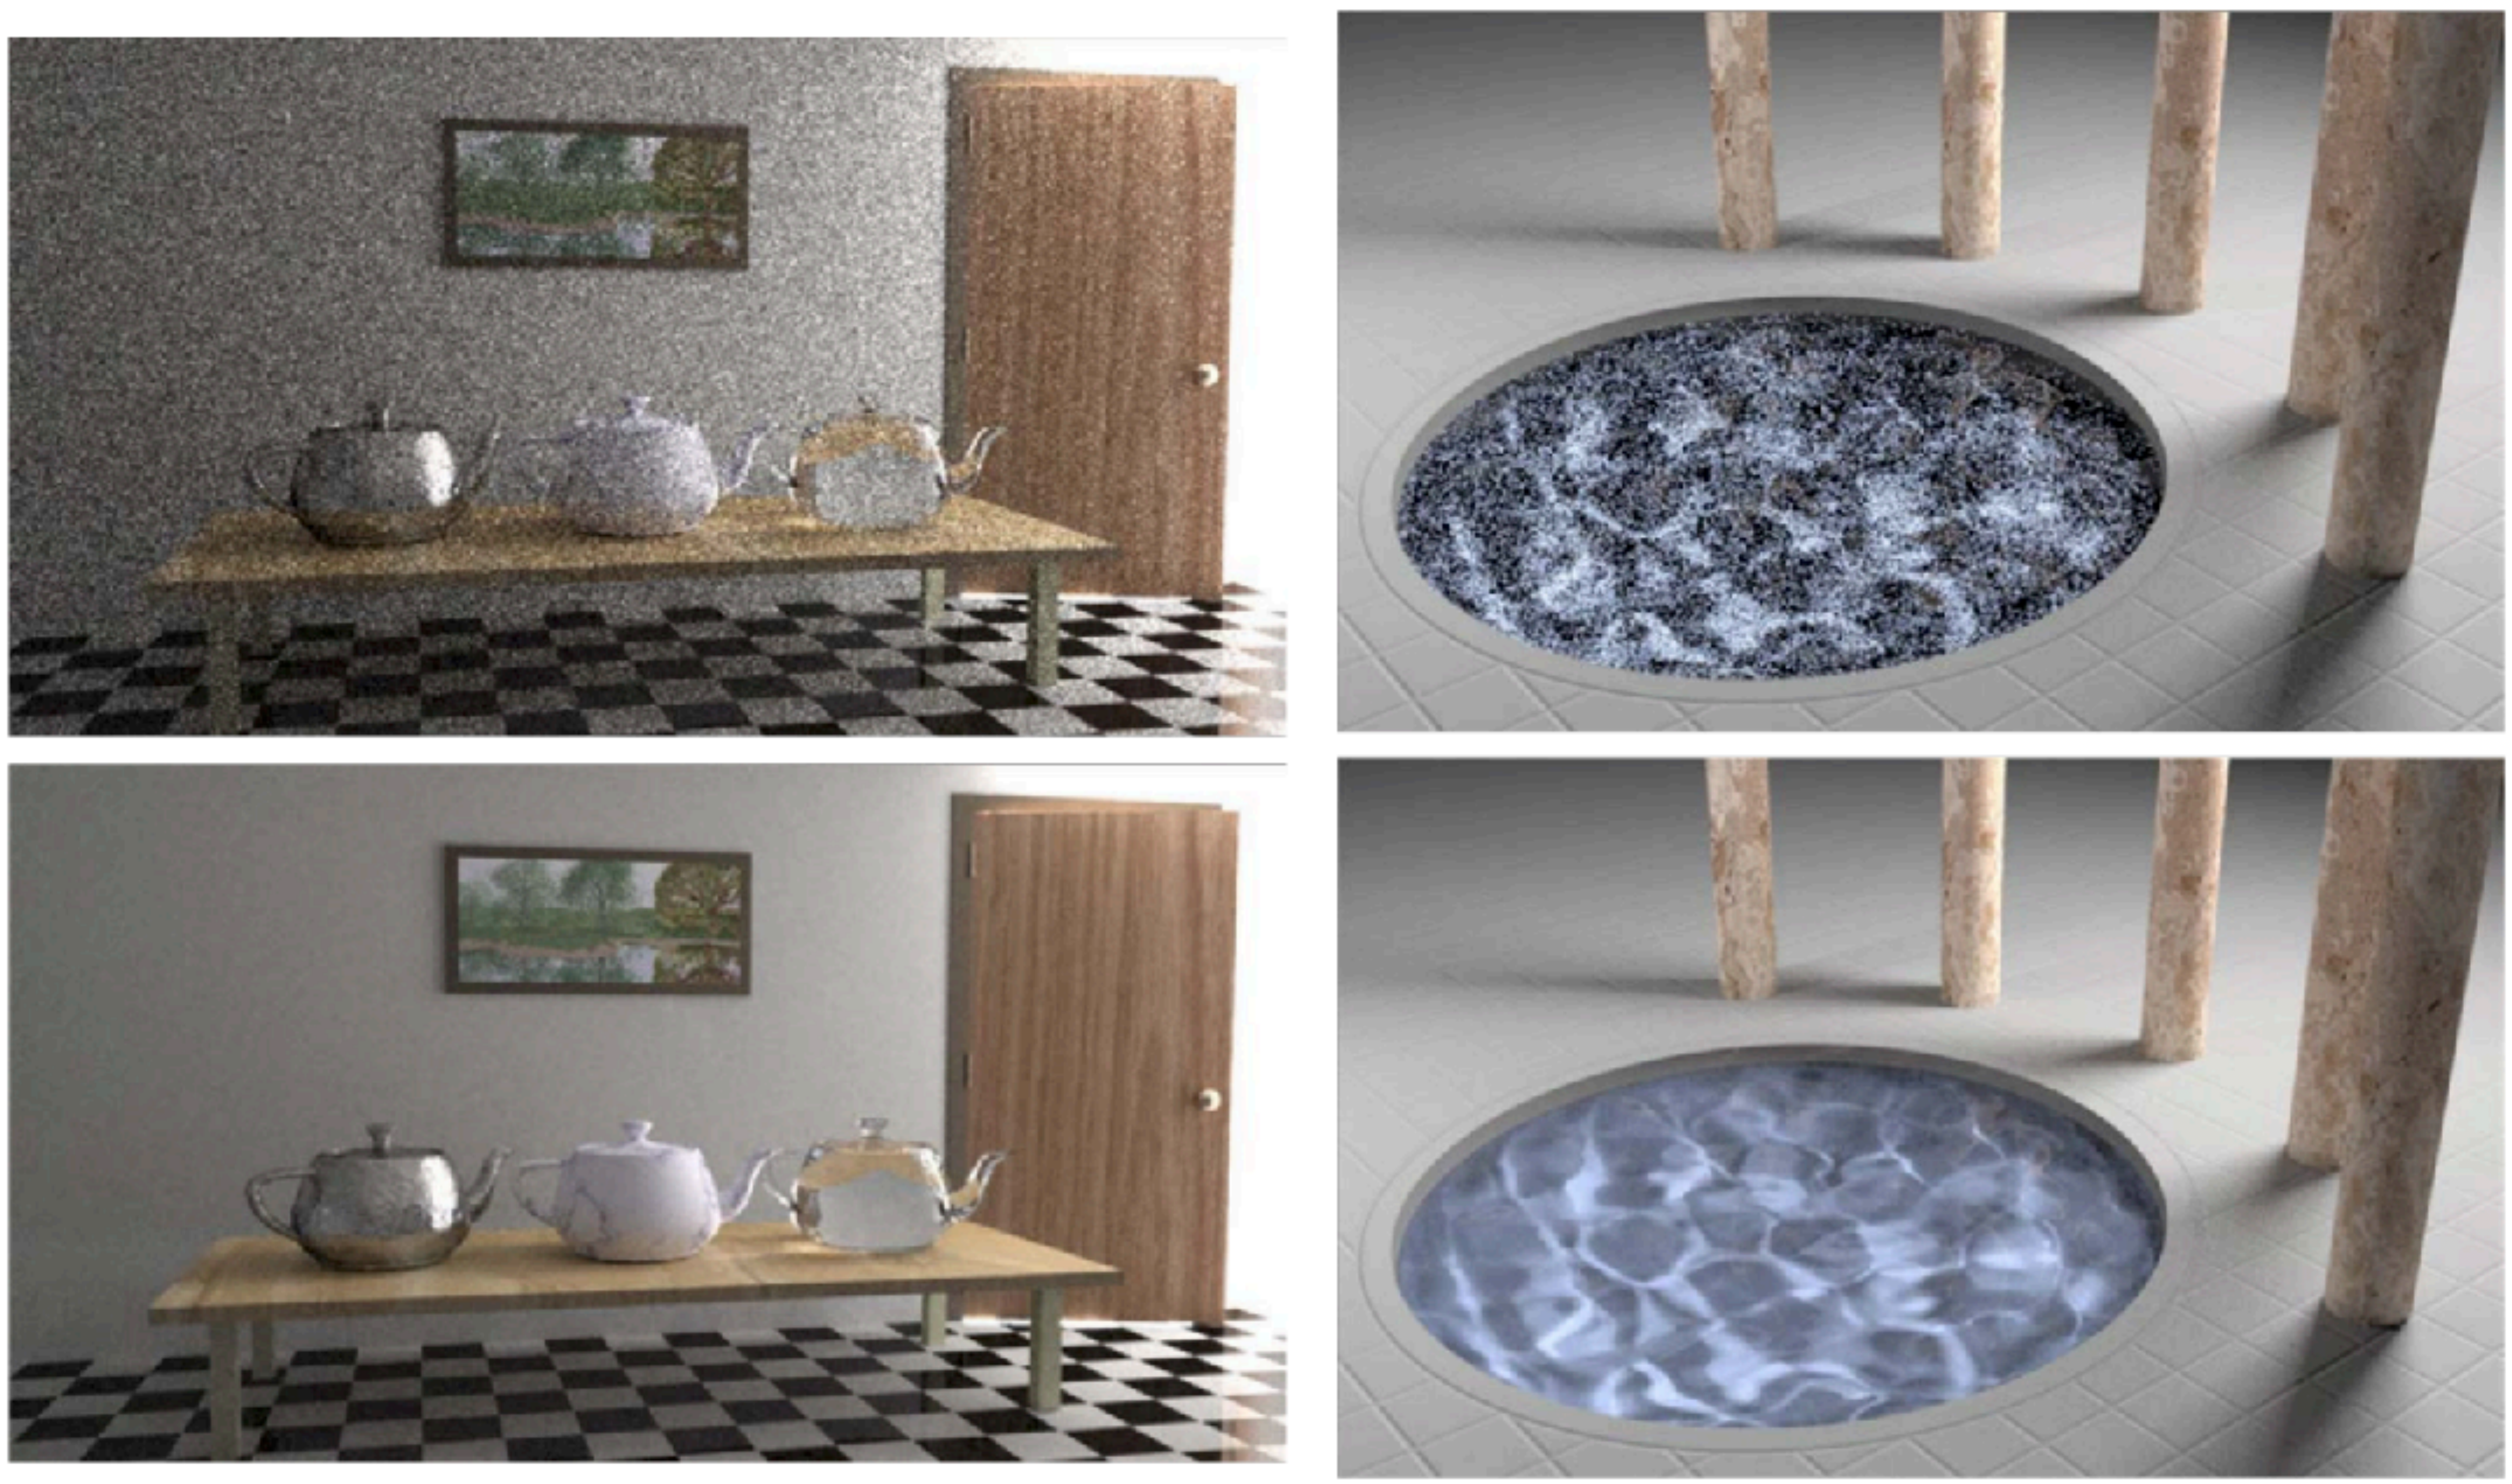

Metropolis Light Transport (MLT) — Cons

- Difficult to estimate the convergence rate
- Does not guarantee equal convergence rate per pixel
- So, usually produces "dirty" results
- Therefore, usually not used to render animations

Photon Mapping

- A biased approach & A two-stage method
- Very good at handling Specular-Diffuse-Specular (SDS) paths and generating caustics

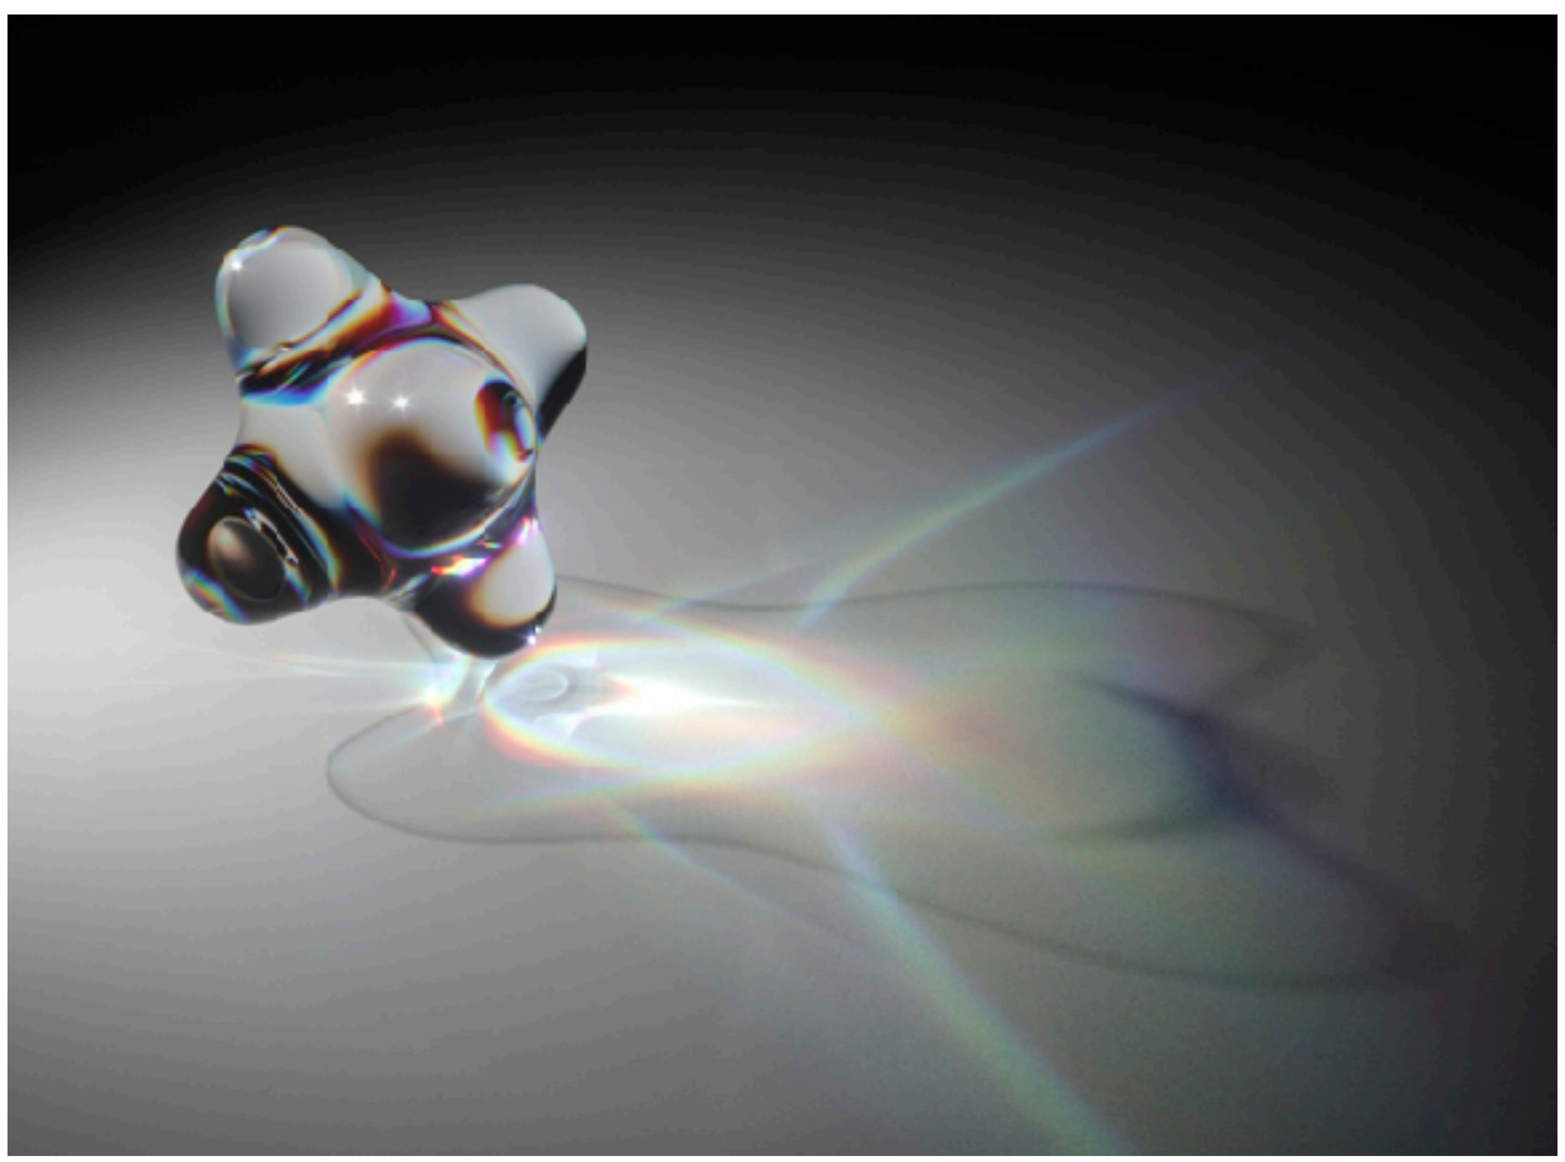

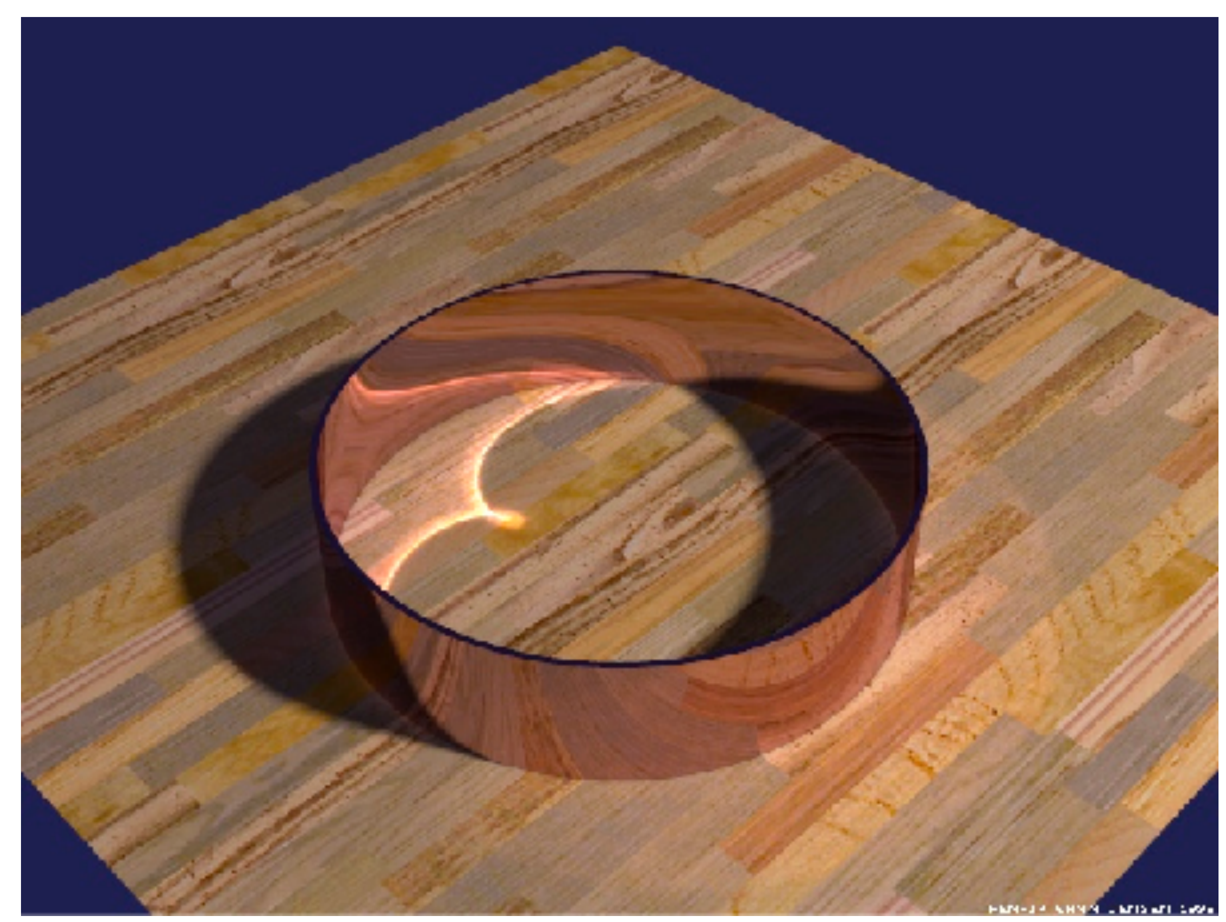

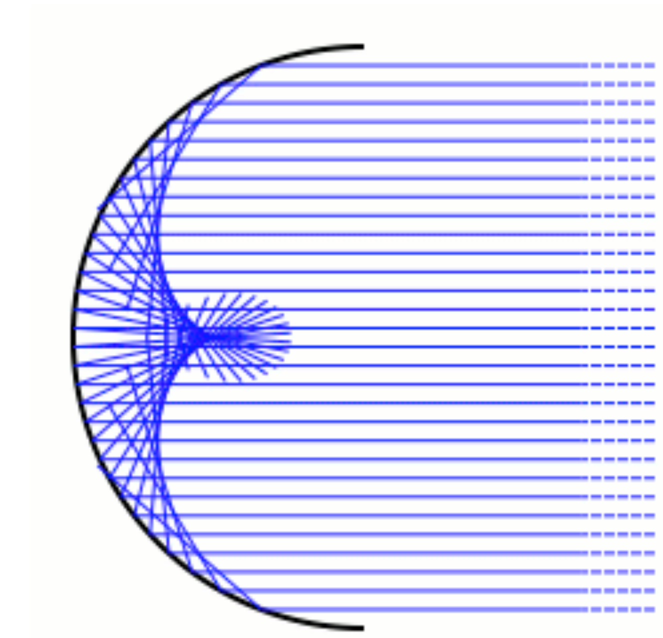

Photon Mapping — Approach (variations apply)

- Stage 1 — photon tracing
  - Emitting photons from the light source, bouncing them around, then recording photons on diffuse surfaces

Stage 1

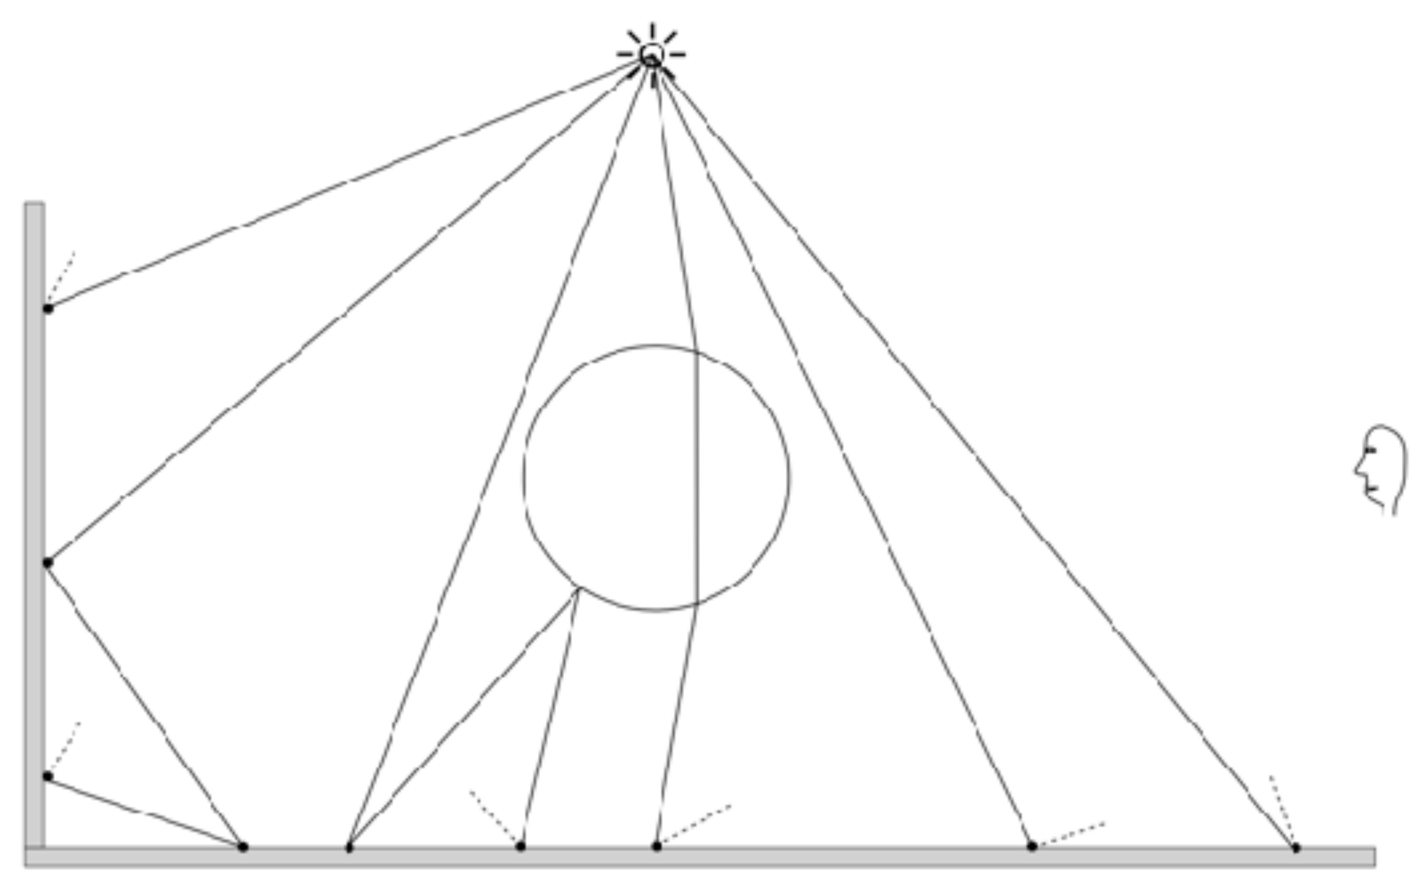

[image courtesy of Henrik Wann Jensen]

- Stage 2 — photon collection (final gathering)
  - Shoot sub-paths from the camera, bouncing them around, until they hit diffuse surfaces

Photon Mapping

- Calculation — local density estimation
  - Idea: areas with more photons should be brighter
  - For each shading point, find the nearest N photons. Take the surface area they over

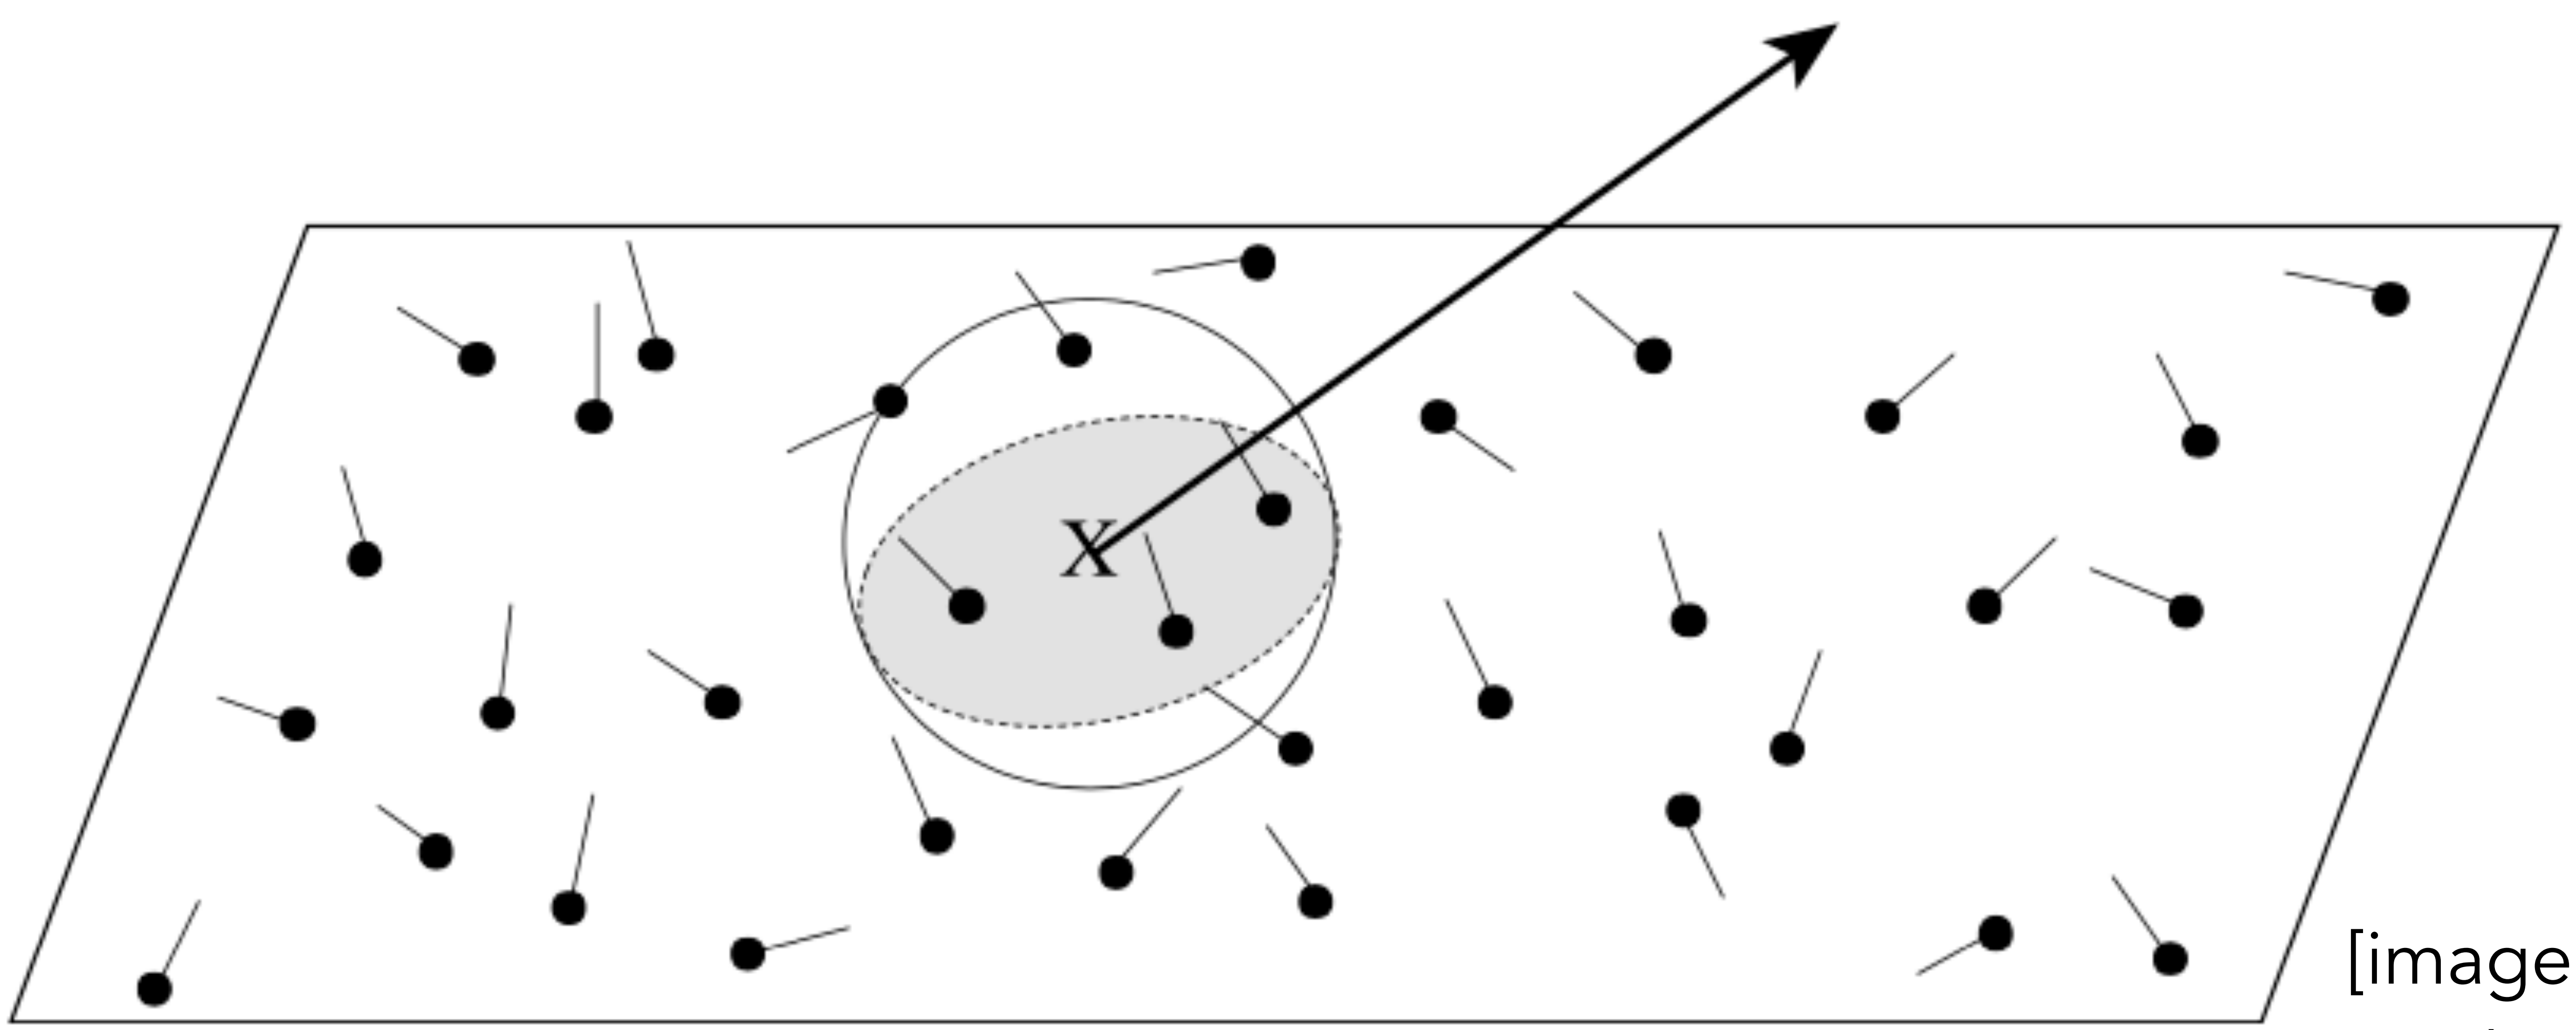

[image courtesy of Henrik Wann Jensen]

Photon Mapping

- Why biased?
- Local Density estimation
  dN / dA != ΔN / ΔA
- But in the sense of limit
  - More photons emitted ->
  - the same N photons
    covers a smaller ΔA ->
  - ΔA is closer to dA
- So, biased but consistent!

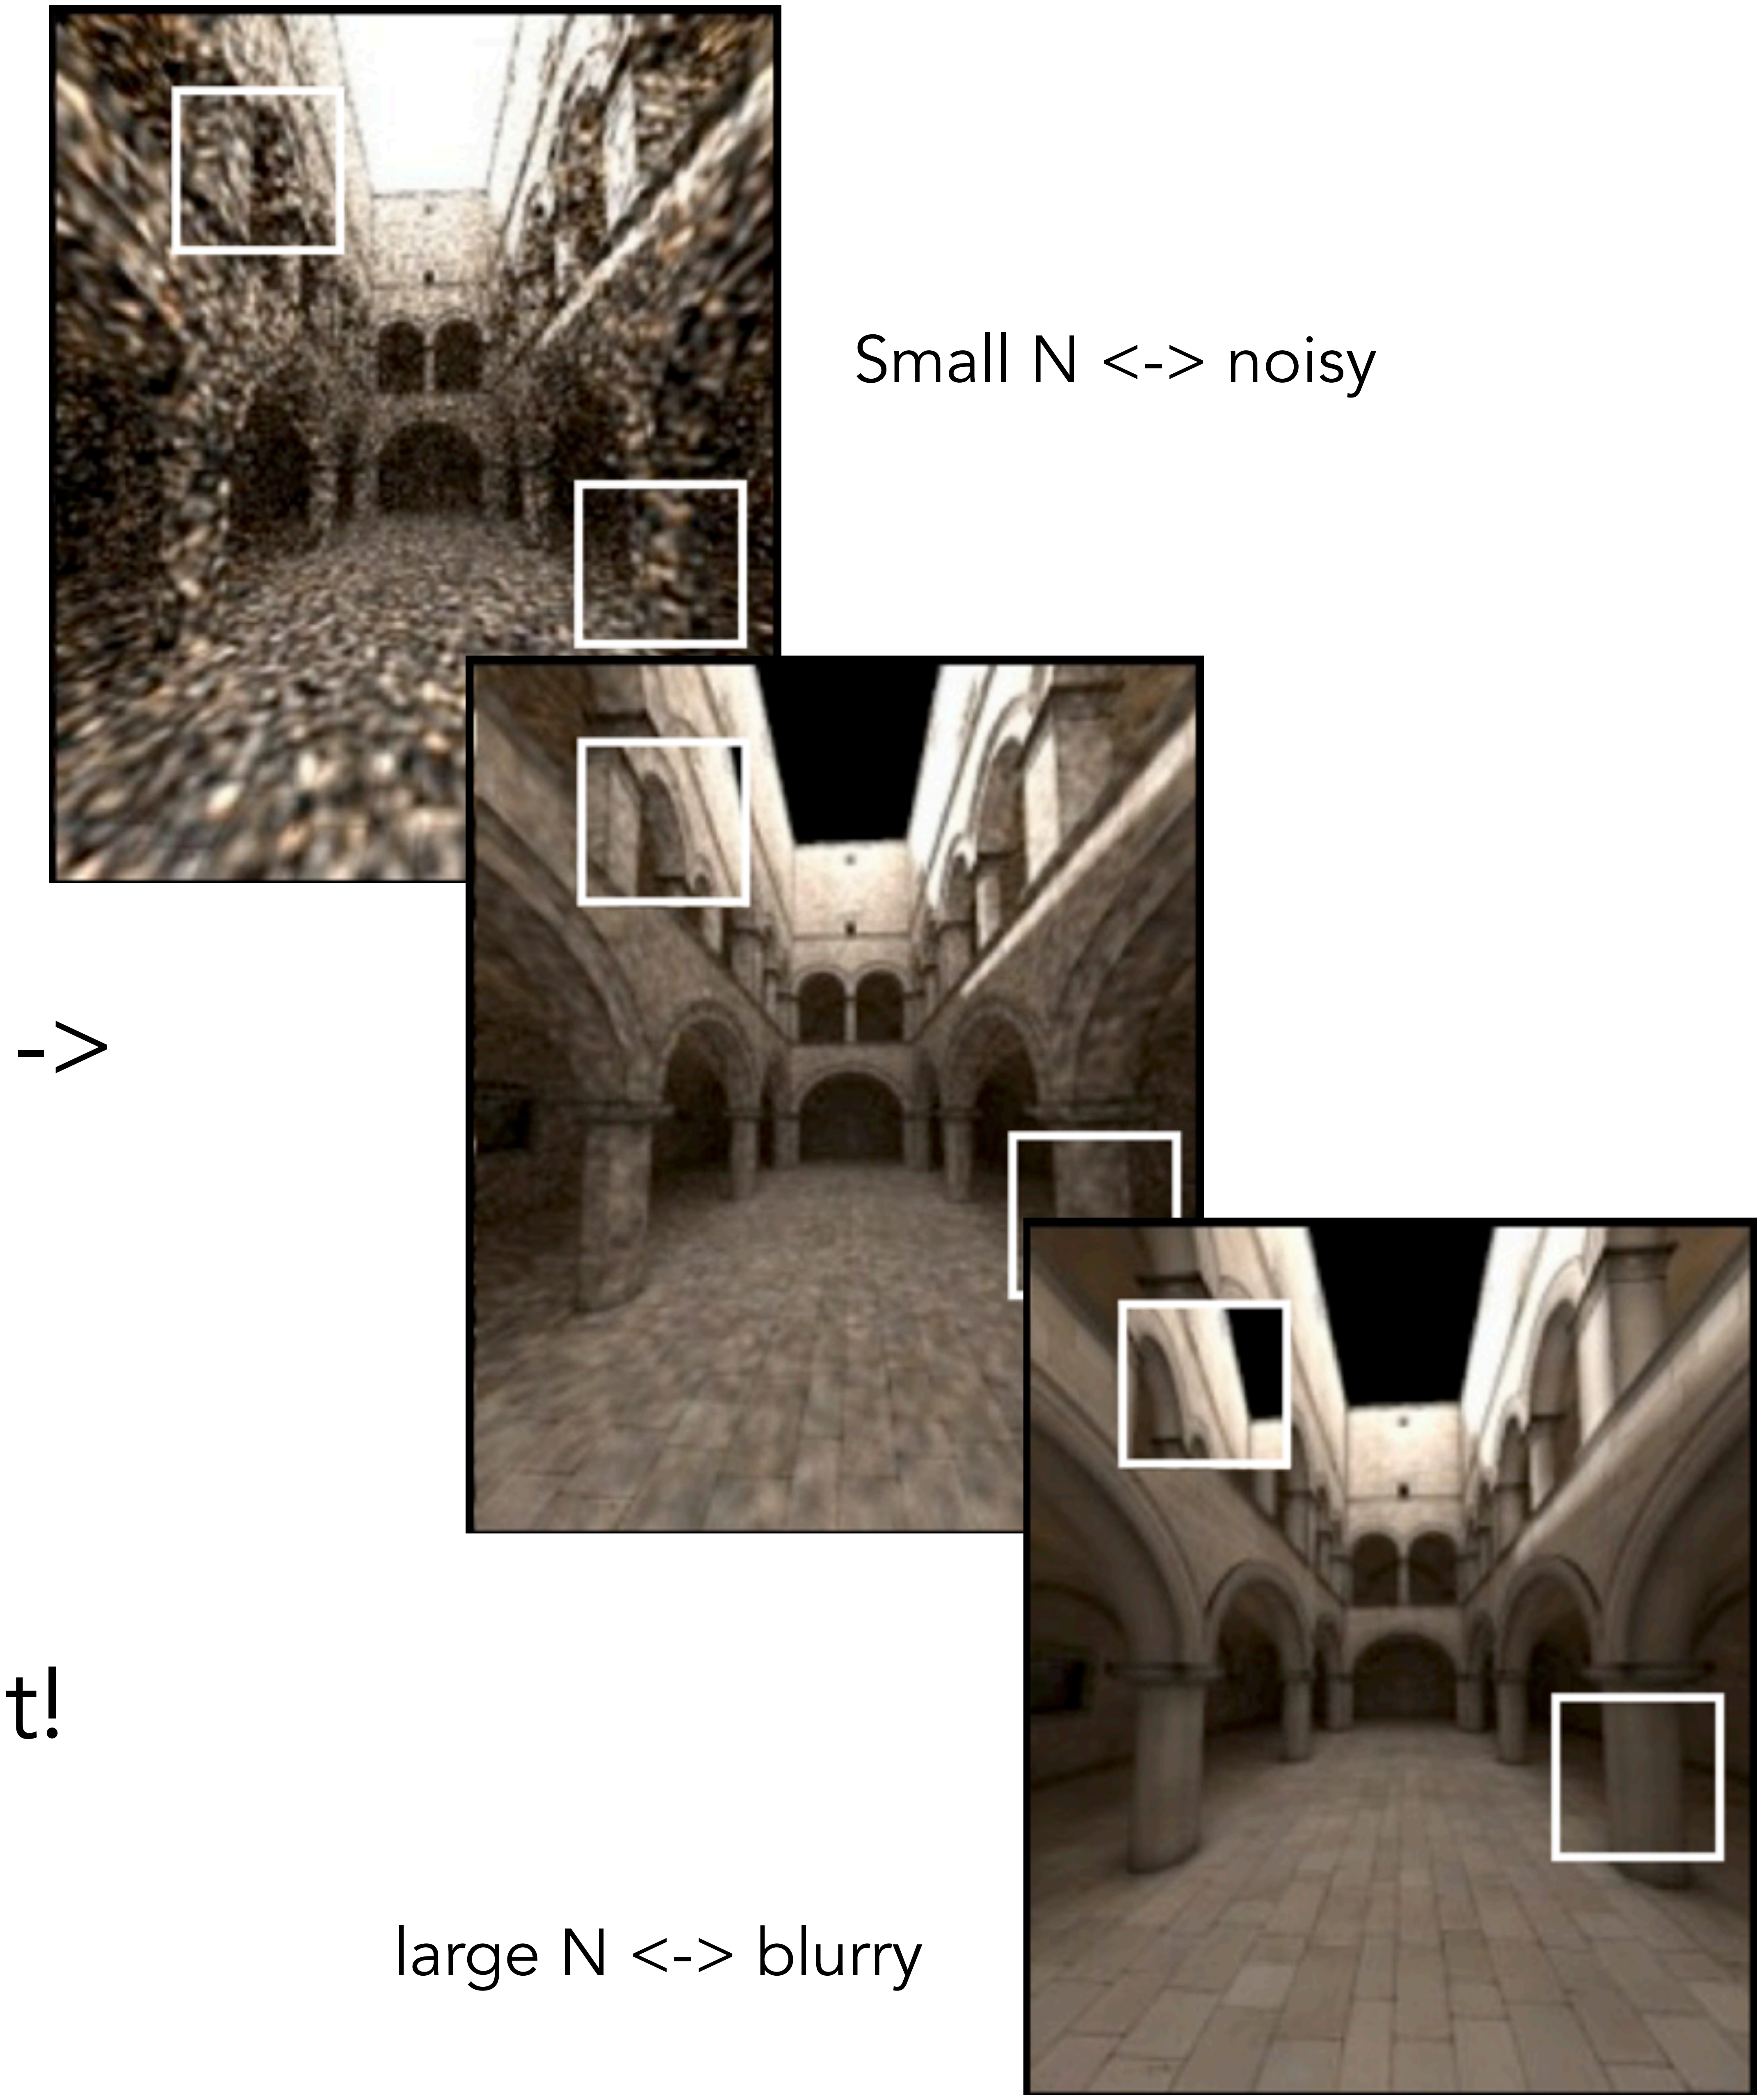

Small N <-> noisy

large N <-> blurry

Photon Mapping

- An easier understanding bias in rendering
  - Biased == blurry
  - Consistent == not blurry with infinite #samples

- Why not do a “const range” search for density estimation?

Vertex Connection and Merging

- A combination of BDPT and Photon Mapping
- Key idea
  - Let's not waste the sub-paths in BDPT if their end points cannot be connected but can be merged
  - Use photon mapping to handle the merging of nearby "photons"

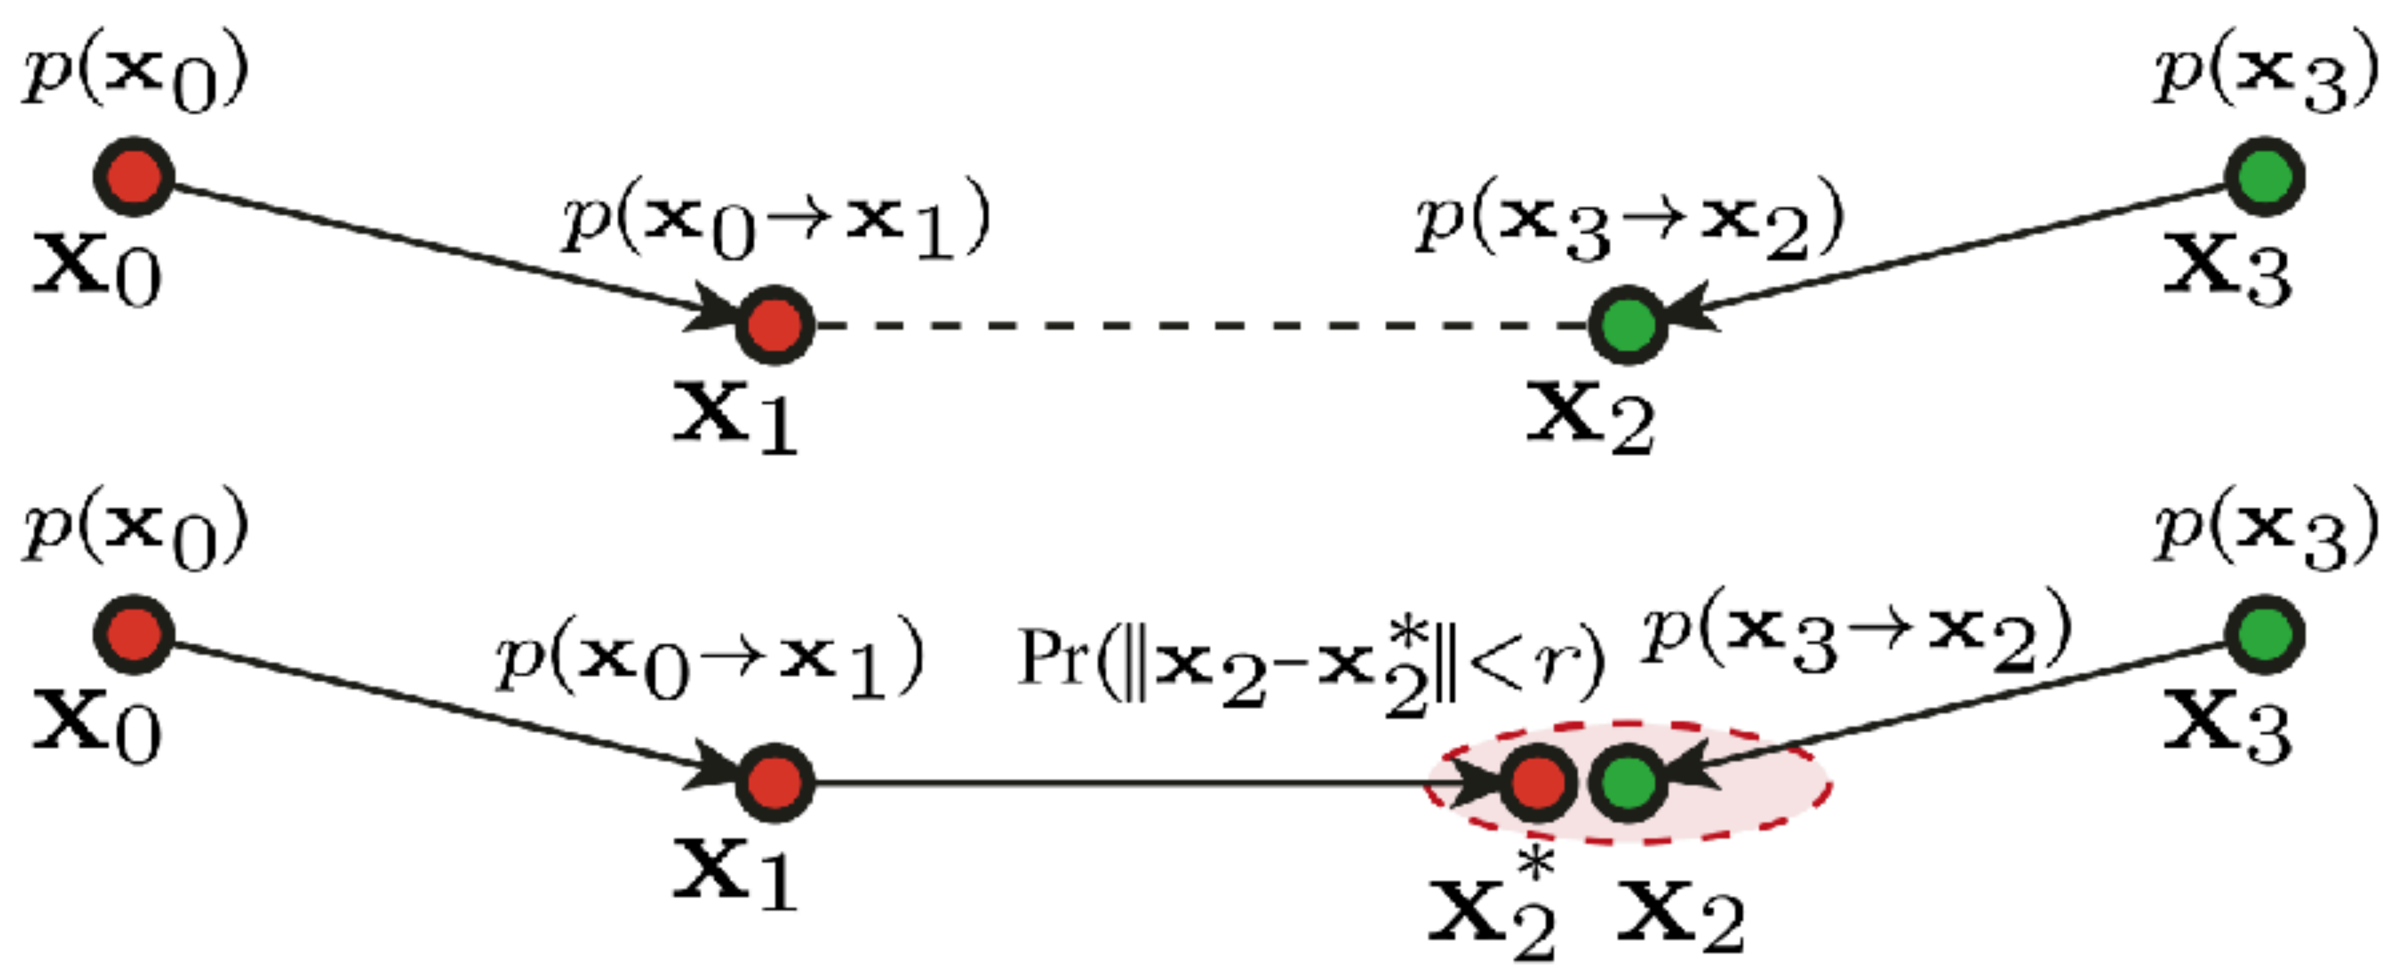

[Georgiev et al. 2012]

Instant Radiosity (IR)

- Sometimes also called many-light approaches
- Key idea
  - Lit surfaces can be treated as light sources
- Approach
  - Shoot light sub-paths and assume the end point of each sub-path is a Virtual Point Light (VPL)
  - Render the scene as usual using these VPLs

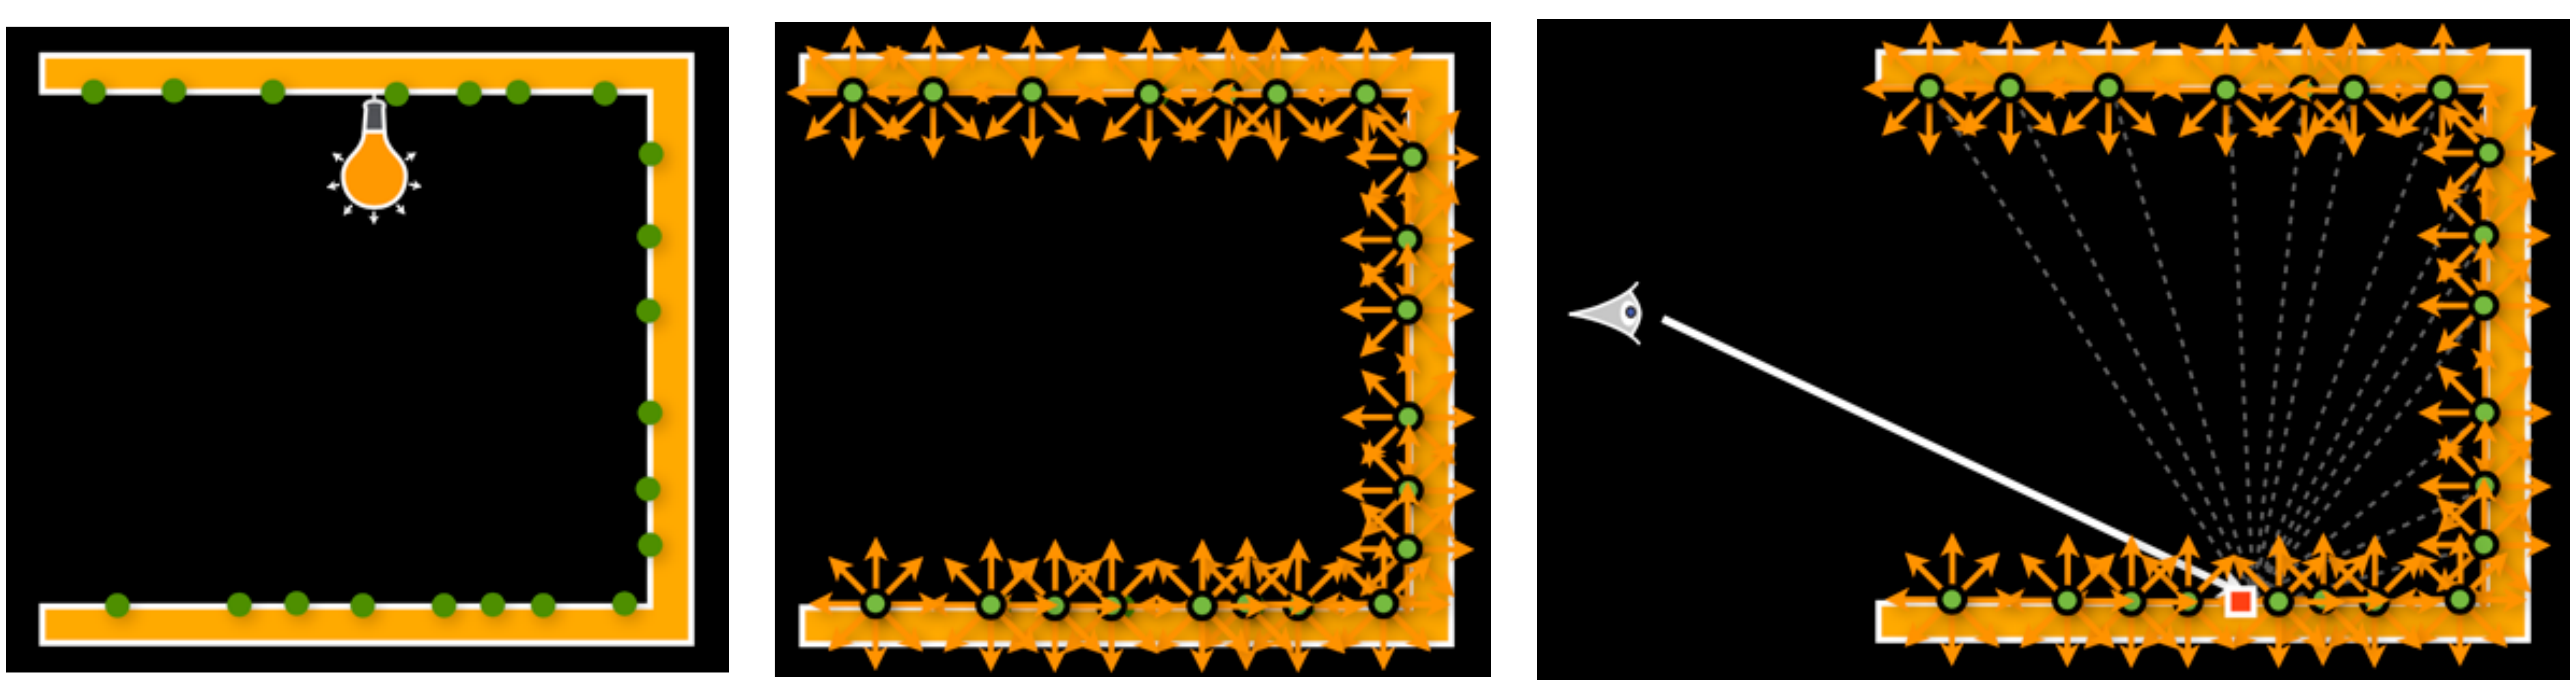

[image courtesy of Derek N.]

Instant Radiosity

- Pros: fast and usually gives good results on diffuse scenes
- Cons
  - Spikes will emerge when VPLs are close to shading points
  - Cannot handle glossy materials

[Liu et al. 2019] (many-light rendering, not IR)

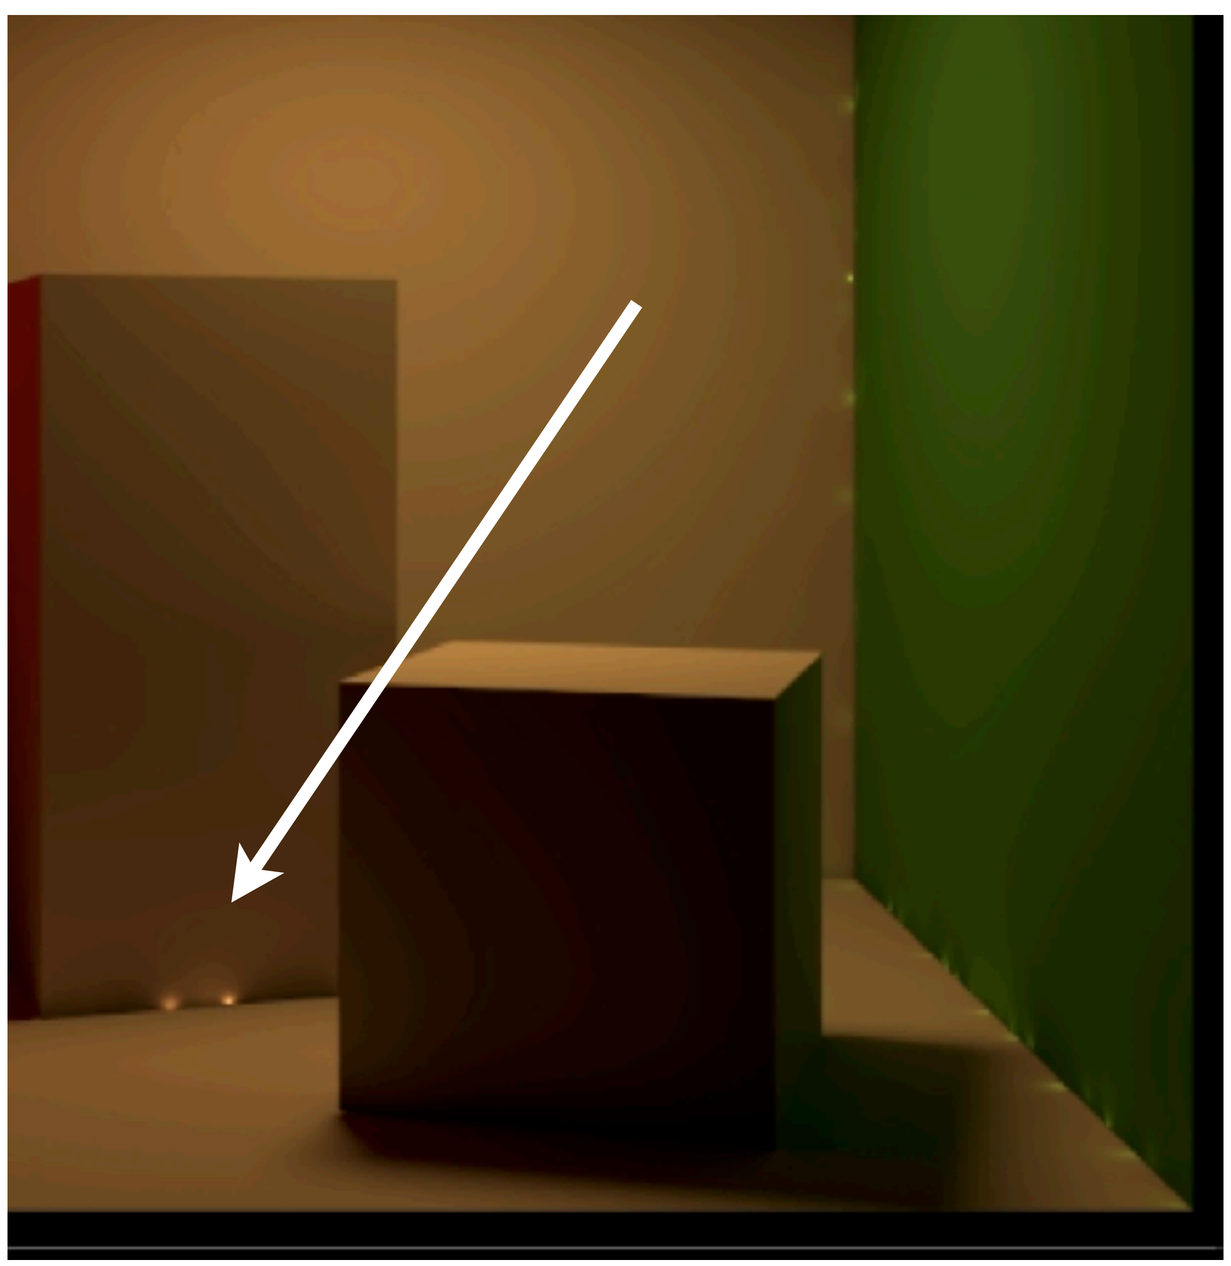

[Rendered using Mitsuba]

Advanced Appearance Modeling

Advanced Appearance Modeling

- Non-surface models
  - Participating media
  - Hair / fur / fiber (BCSDF)
  - Granular material
- Surface models
  - Translucent material (BSSRDF)
  - Cloth
  - Detailed material (non-statistical BRDF)
- Procedural appearance

Non-Surface Models

Participating Media: Fog

[Novák et al. 2012]

Participating Media: Cloud

[by thephotographer0]

Participating Media

- At any point as light travels through a participating medium, it can be (partially) absorbed and scattered.

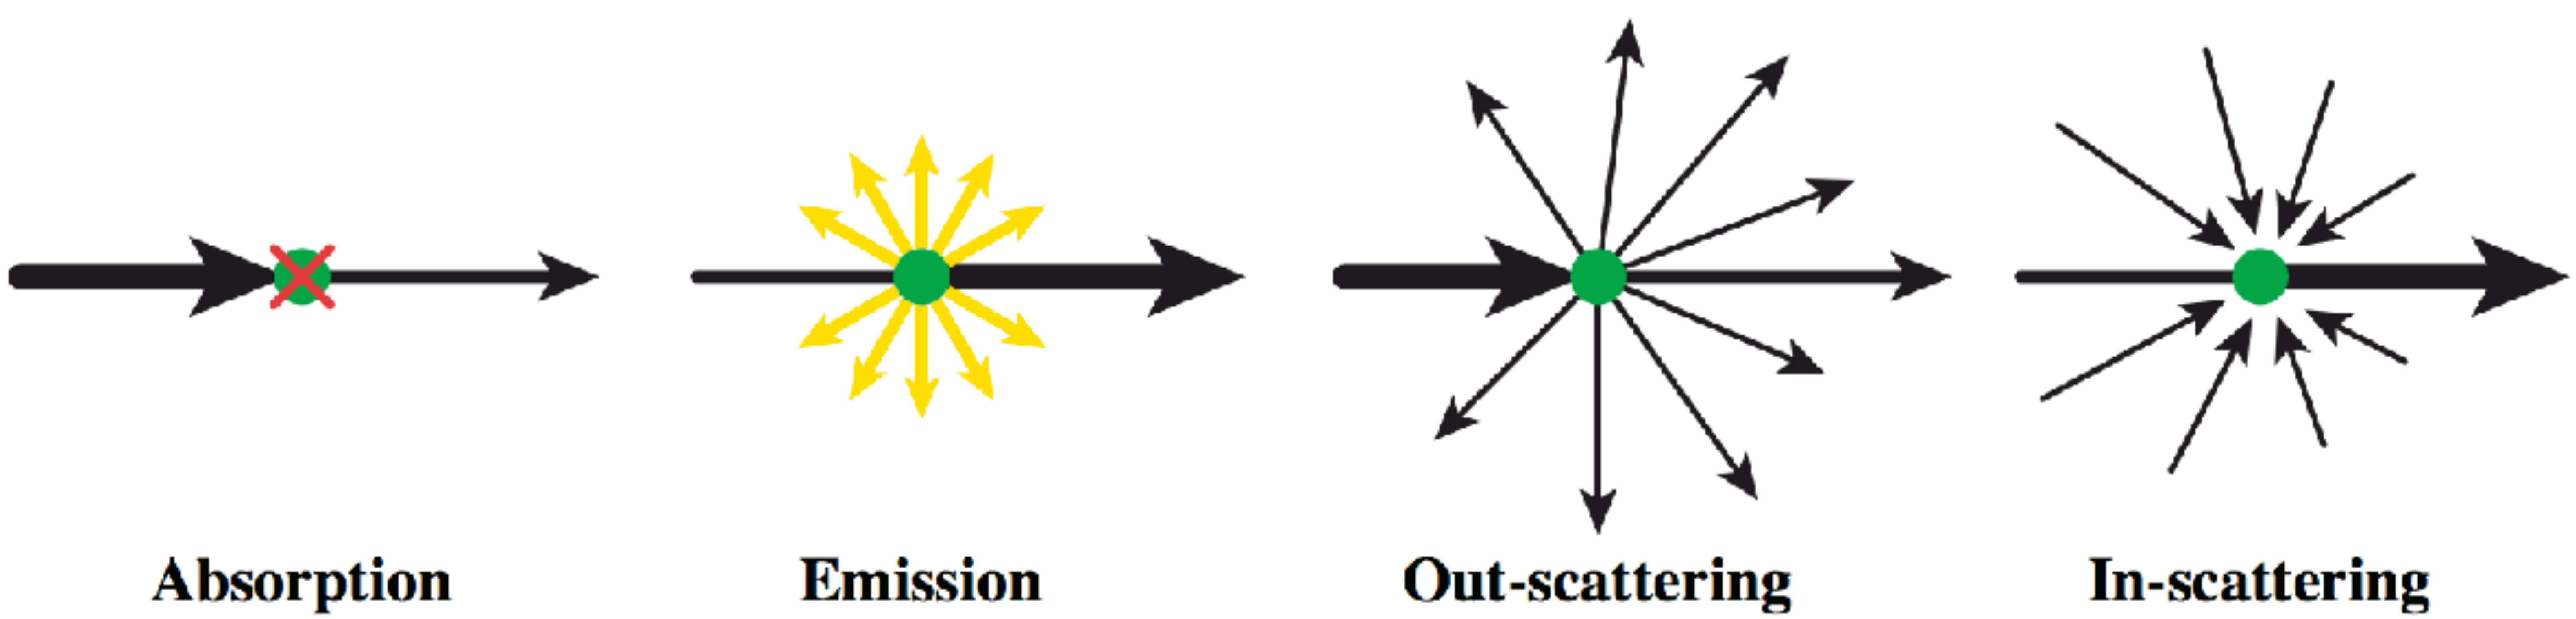

[Wojciech Jarosz]

Participating Media

- Use Phase Function to describe the angular distribution of light scattering at any point x within participating media.

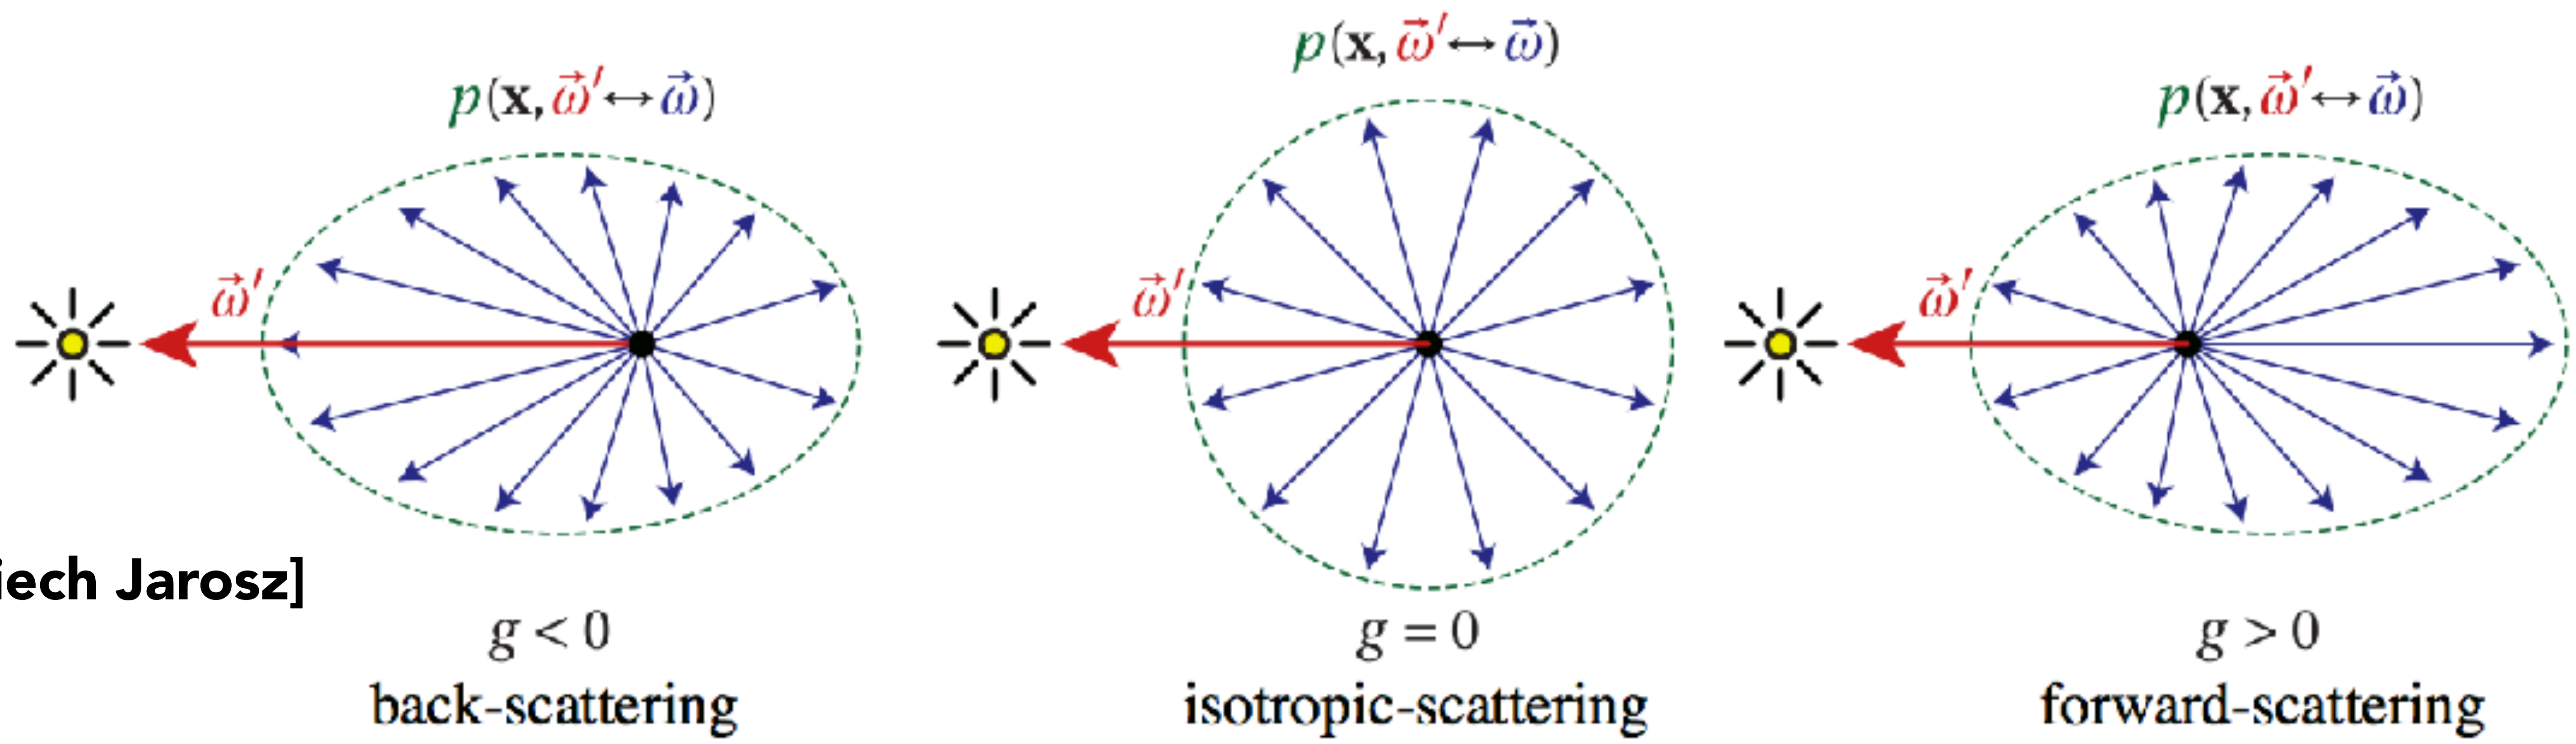

[Wojciech Jarosz]

Participating Media: Rendering

- Randomly choose a direction to bounce
- Randomly choose a distance to go straight
- At each 'shading point', connect to the light

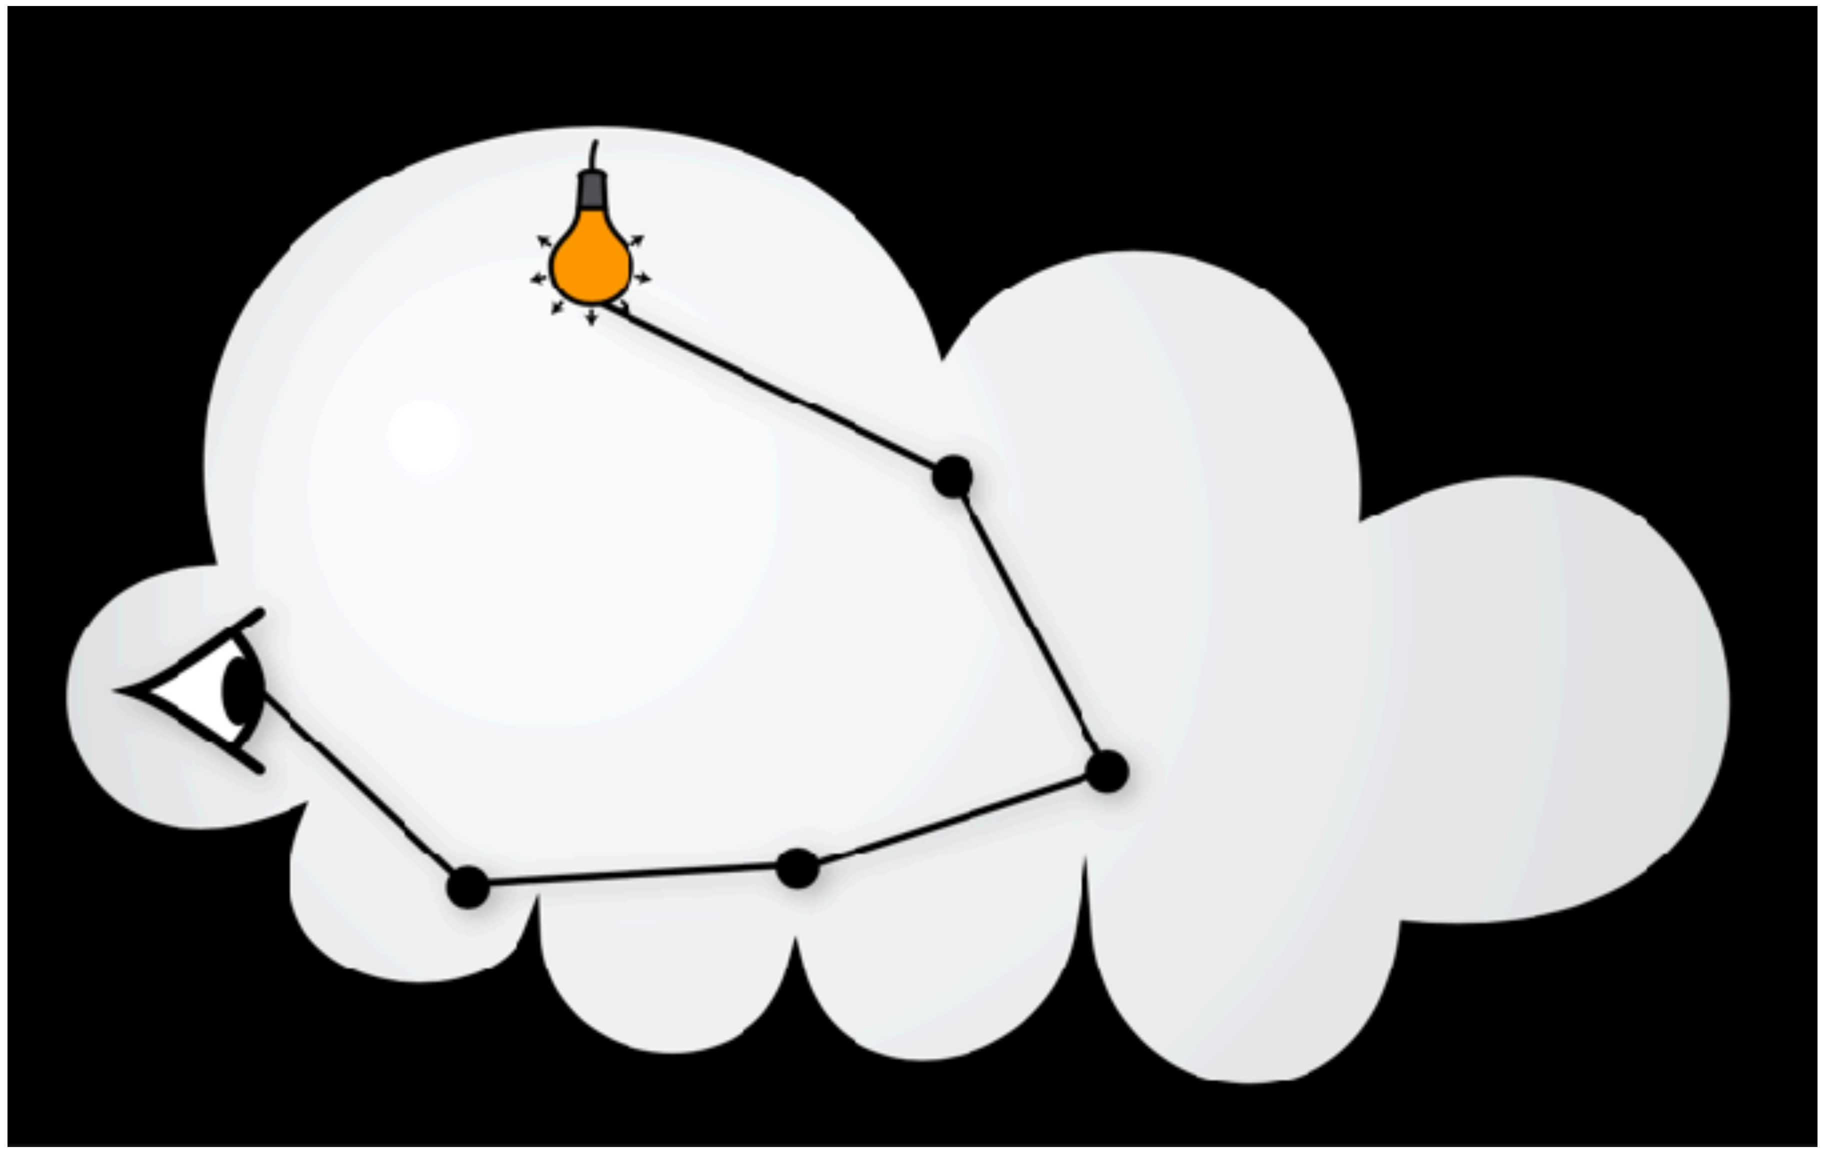

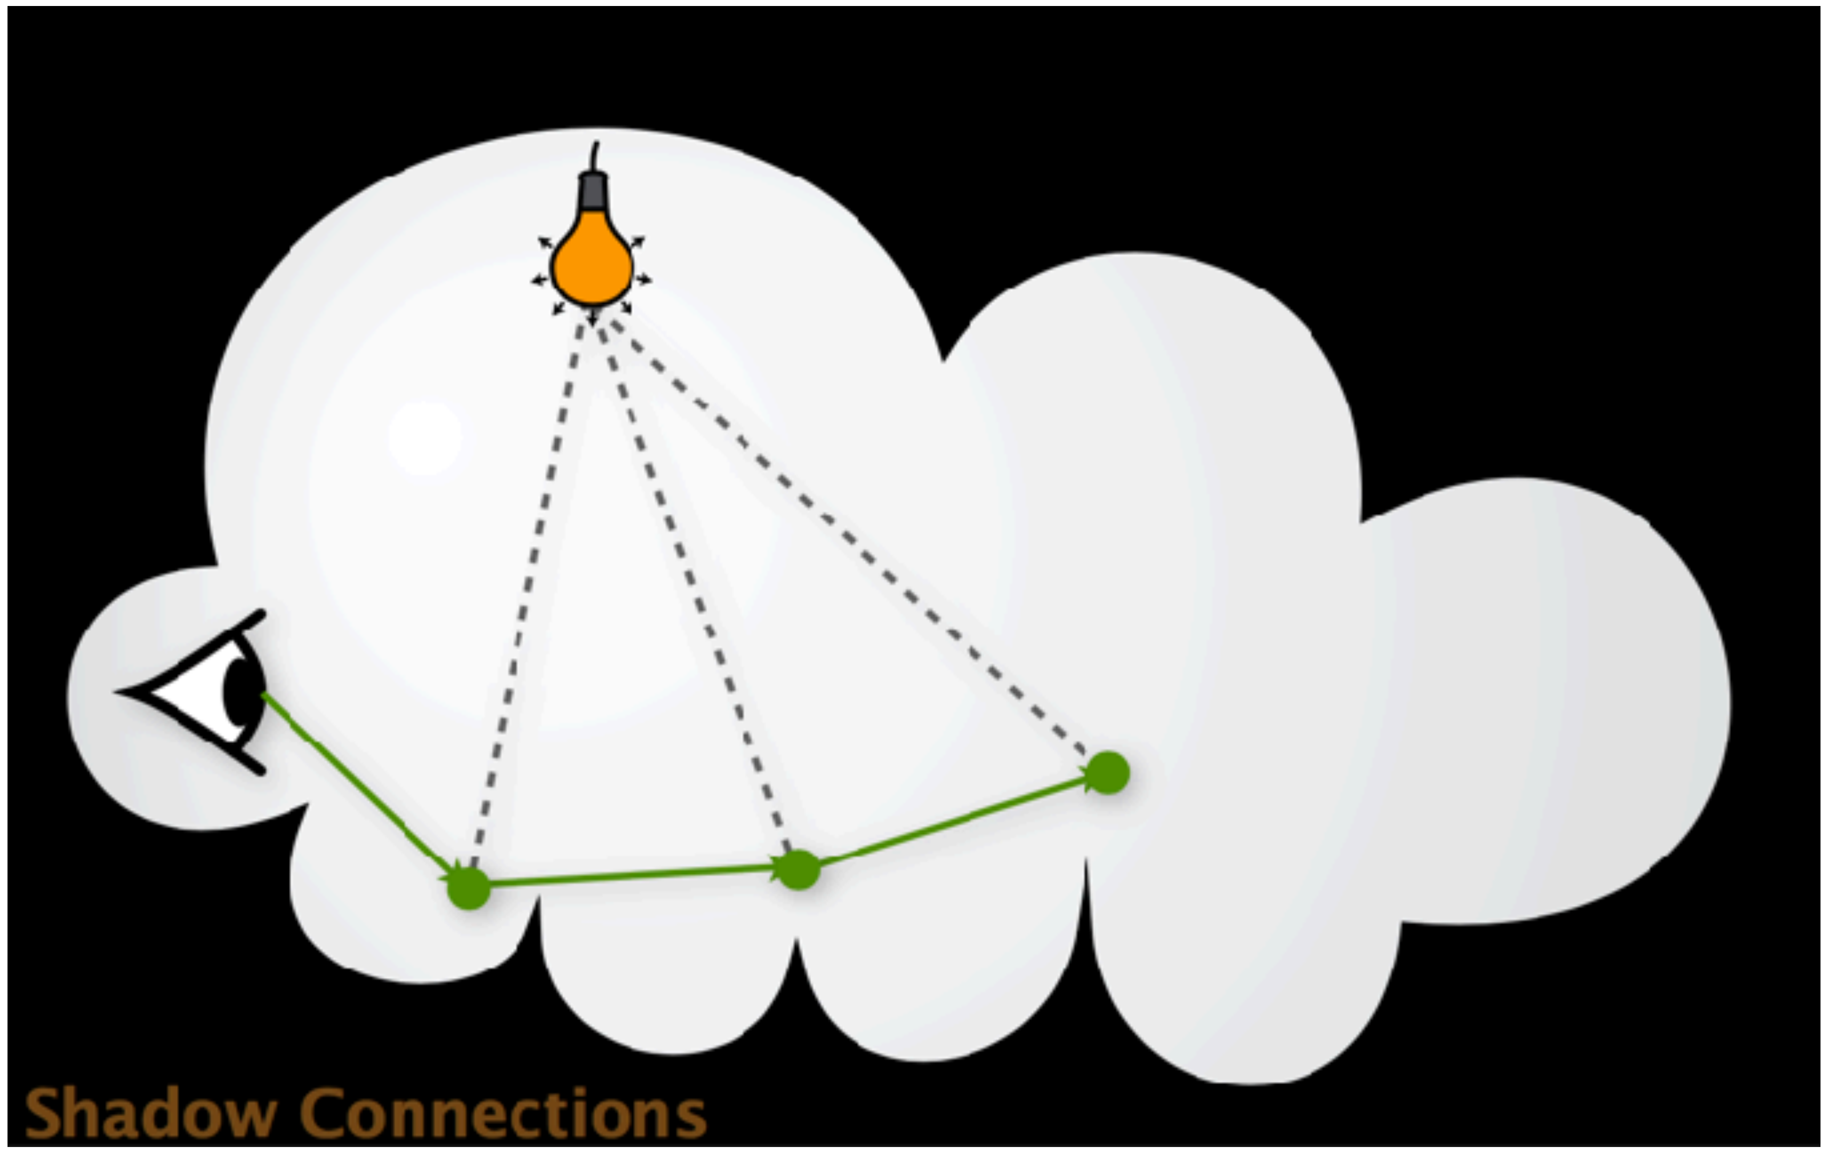

[Derek Nowrouzezahrai]

Participating Media: Application

[Big Hero 6, 2014 Disney]

Participating Media: Application

[Assassin's Creed Syndicate. 2015 Ubisoft]

Participating Media: Demo

[Stomakhin et al. 2014]

Hair Appearance

Kajiya-Kay Model

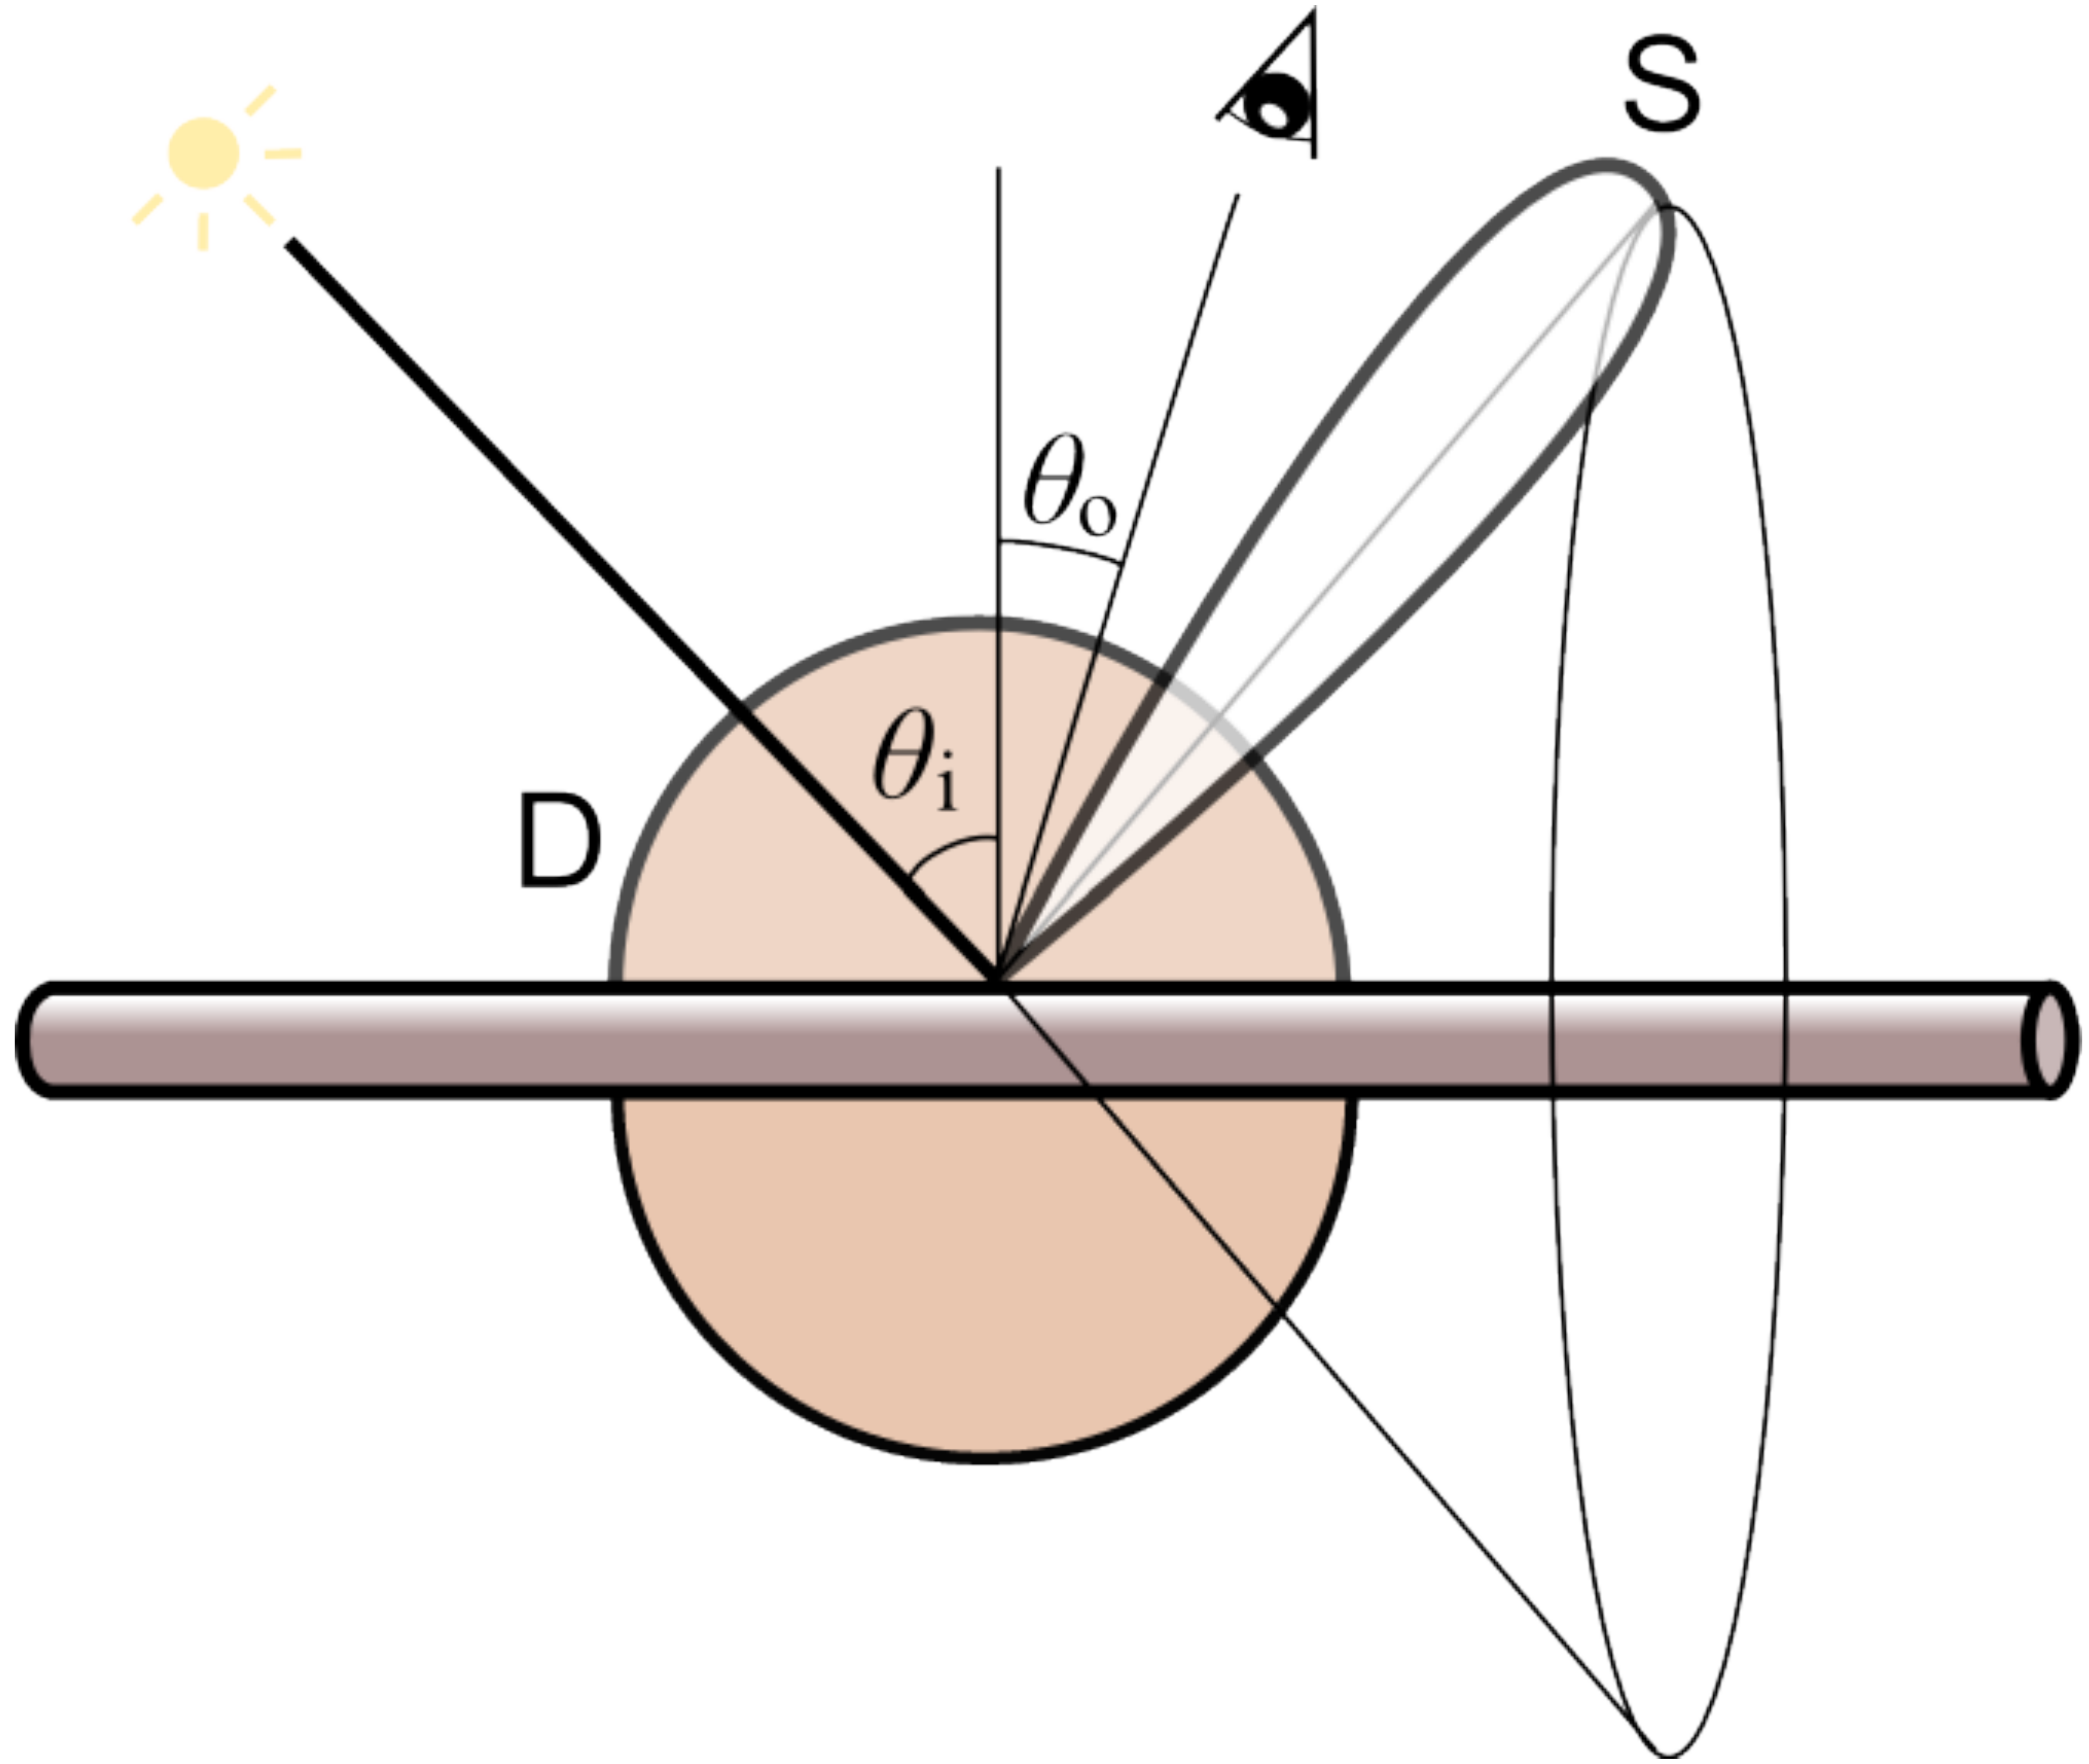

[Image courtesy of Chiwei Tseng]

Kajiya-Kay Model

[Yuksel et al. 2008]

Marschner Model

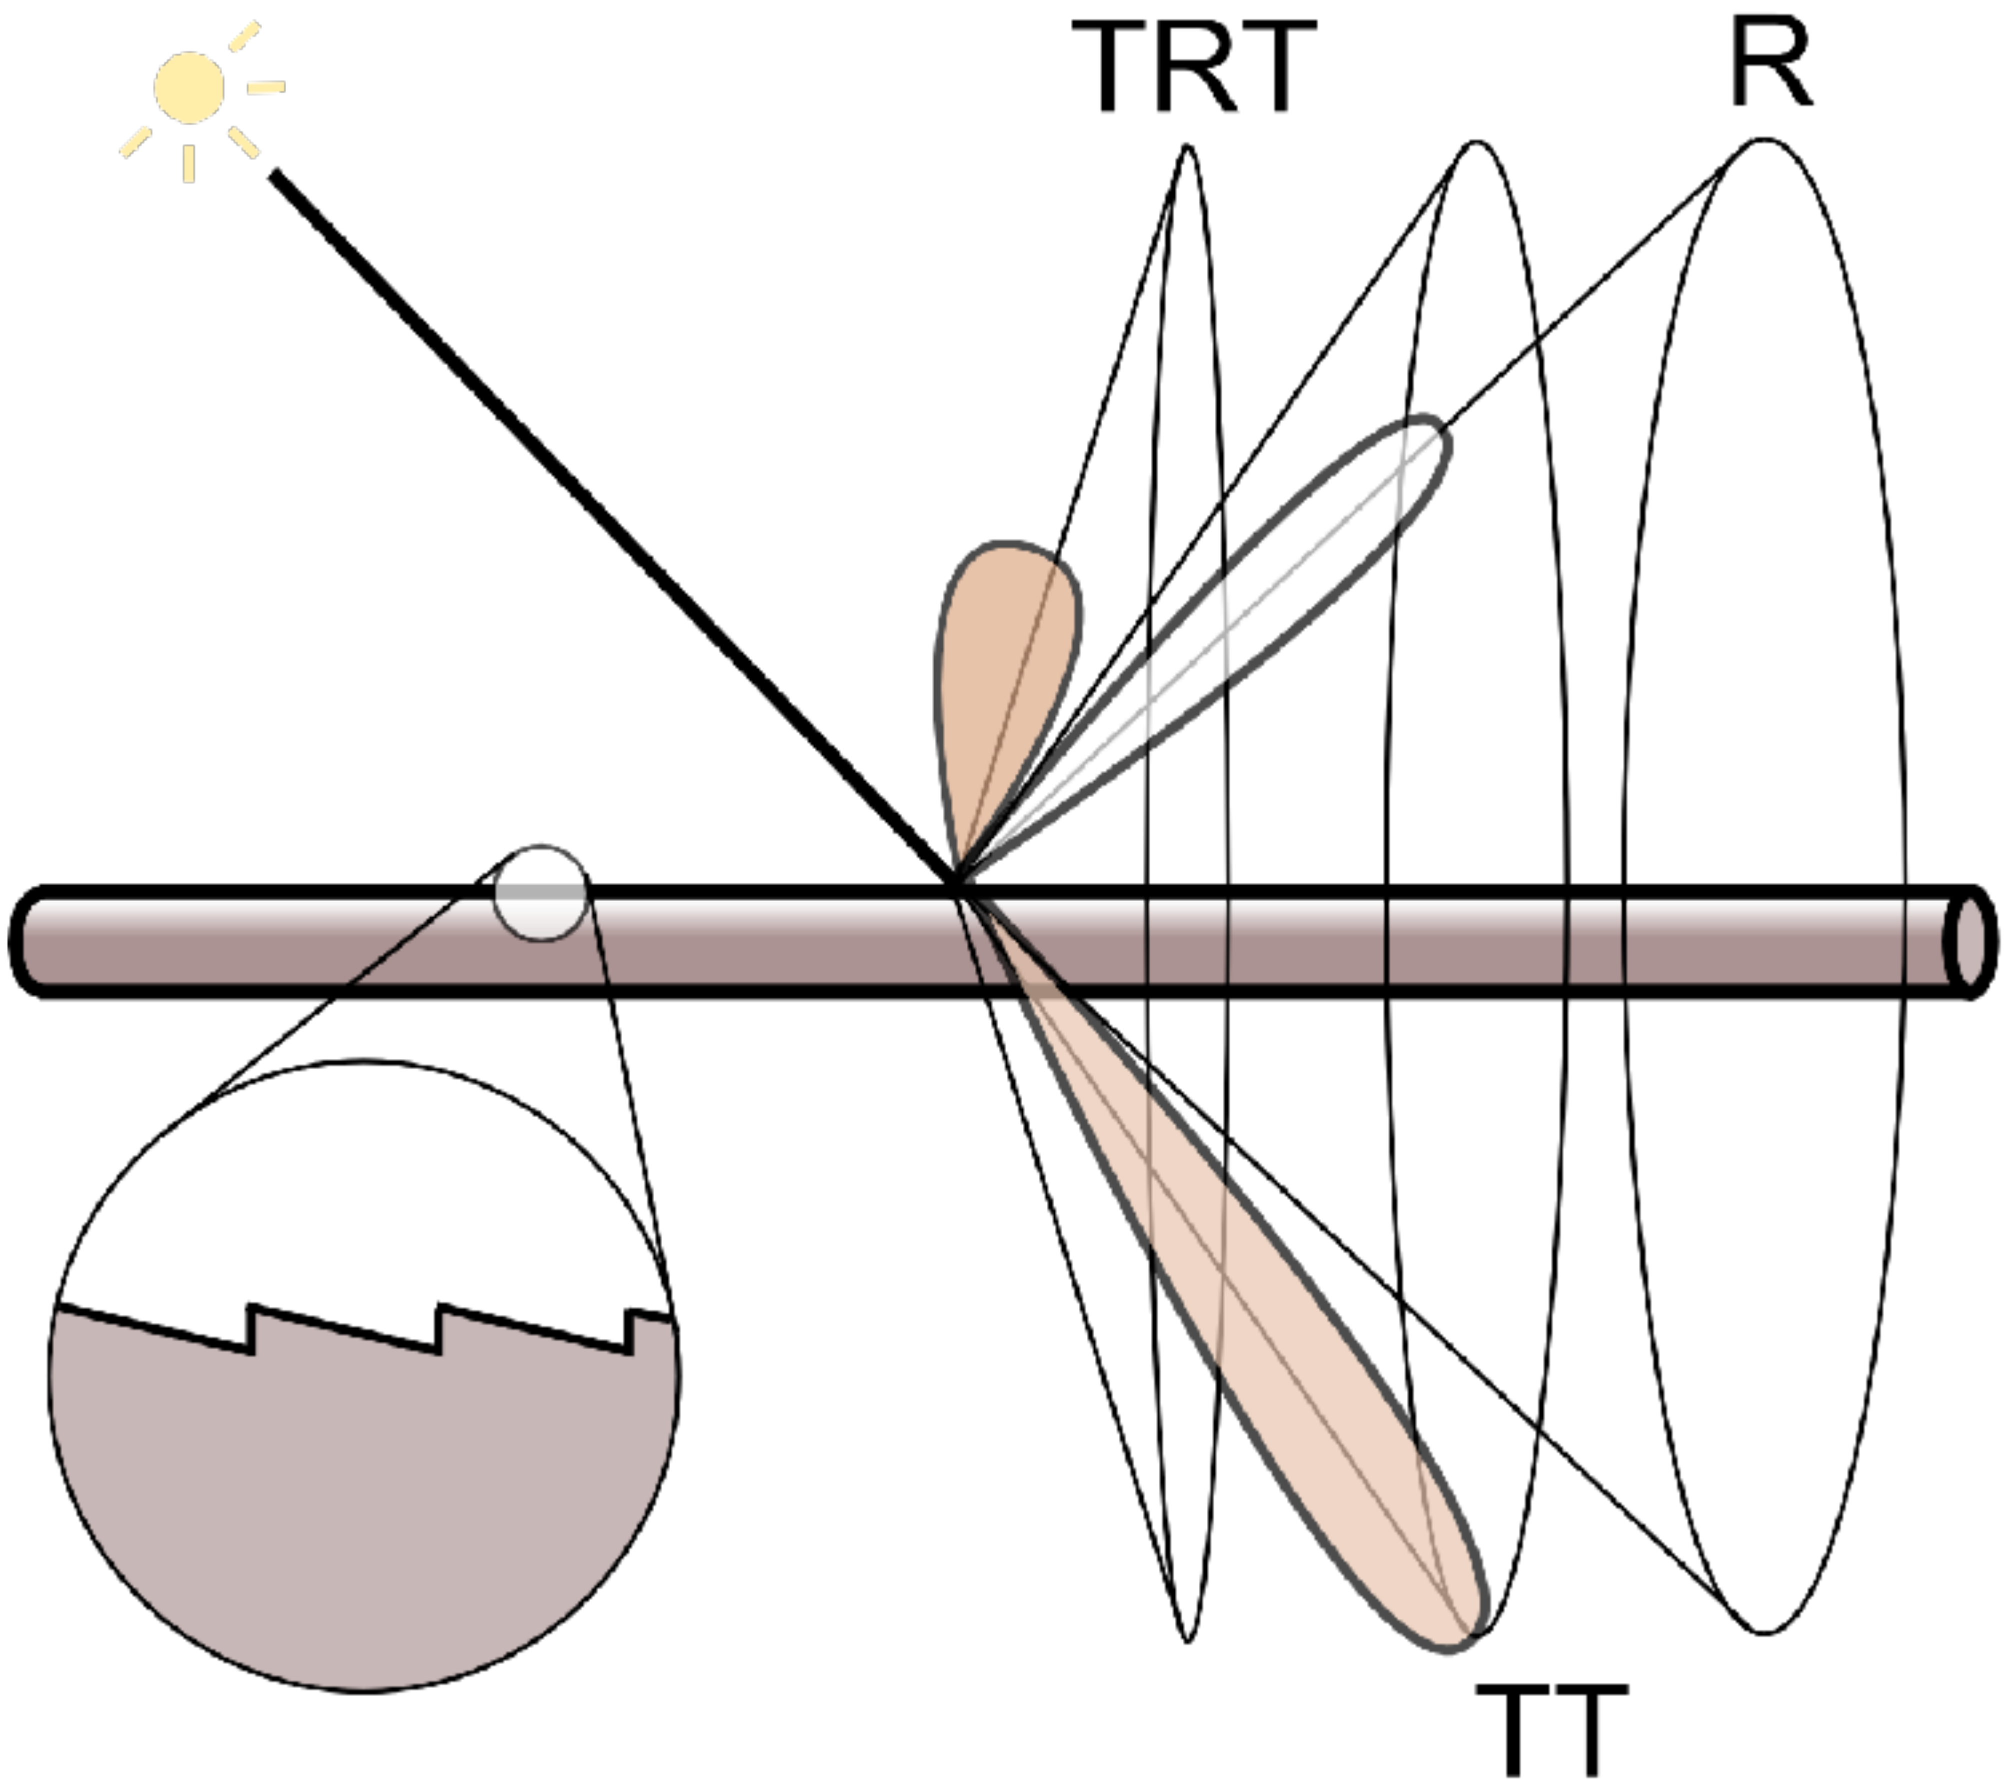

[Image courtesy of Chiwei Tseng]

Marschner Model

- Glass-like cylinder

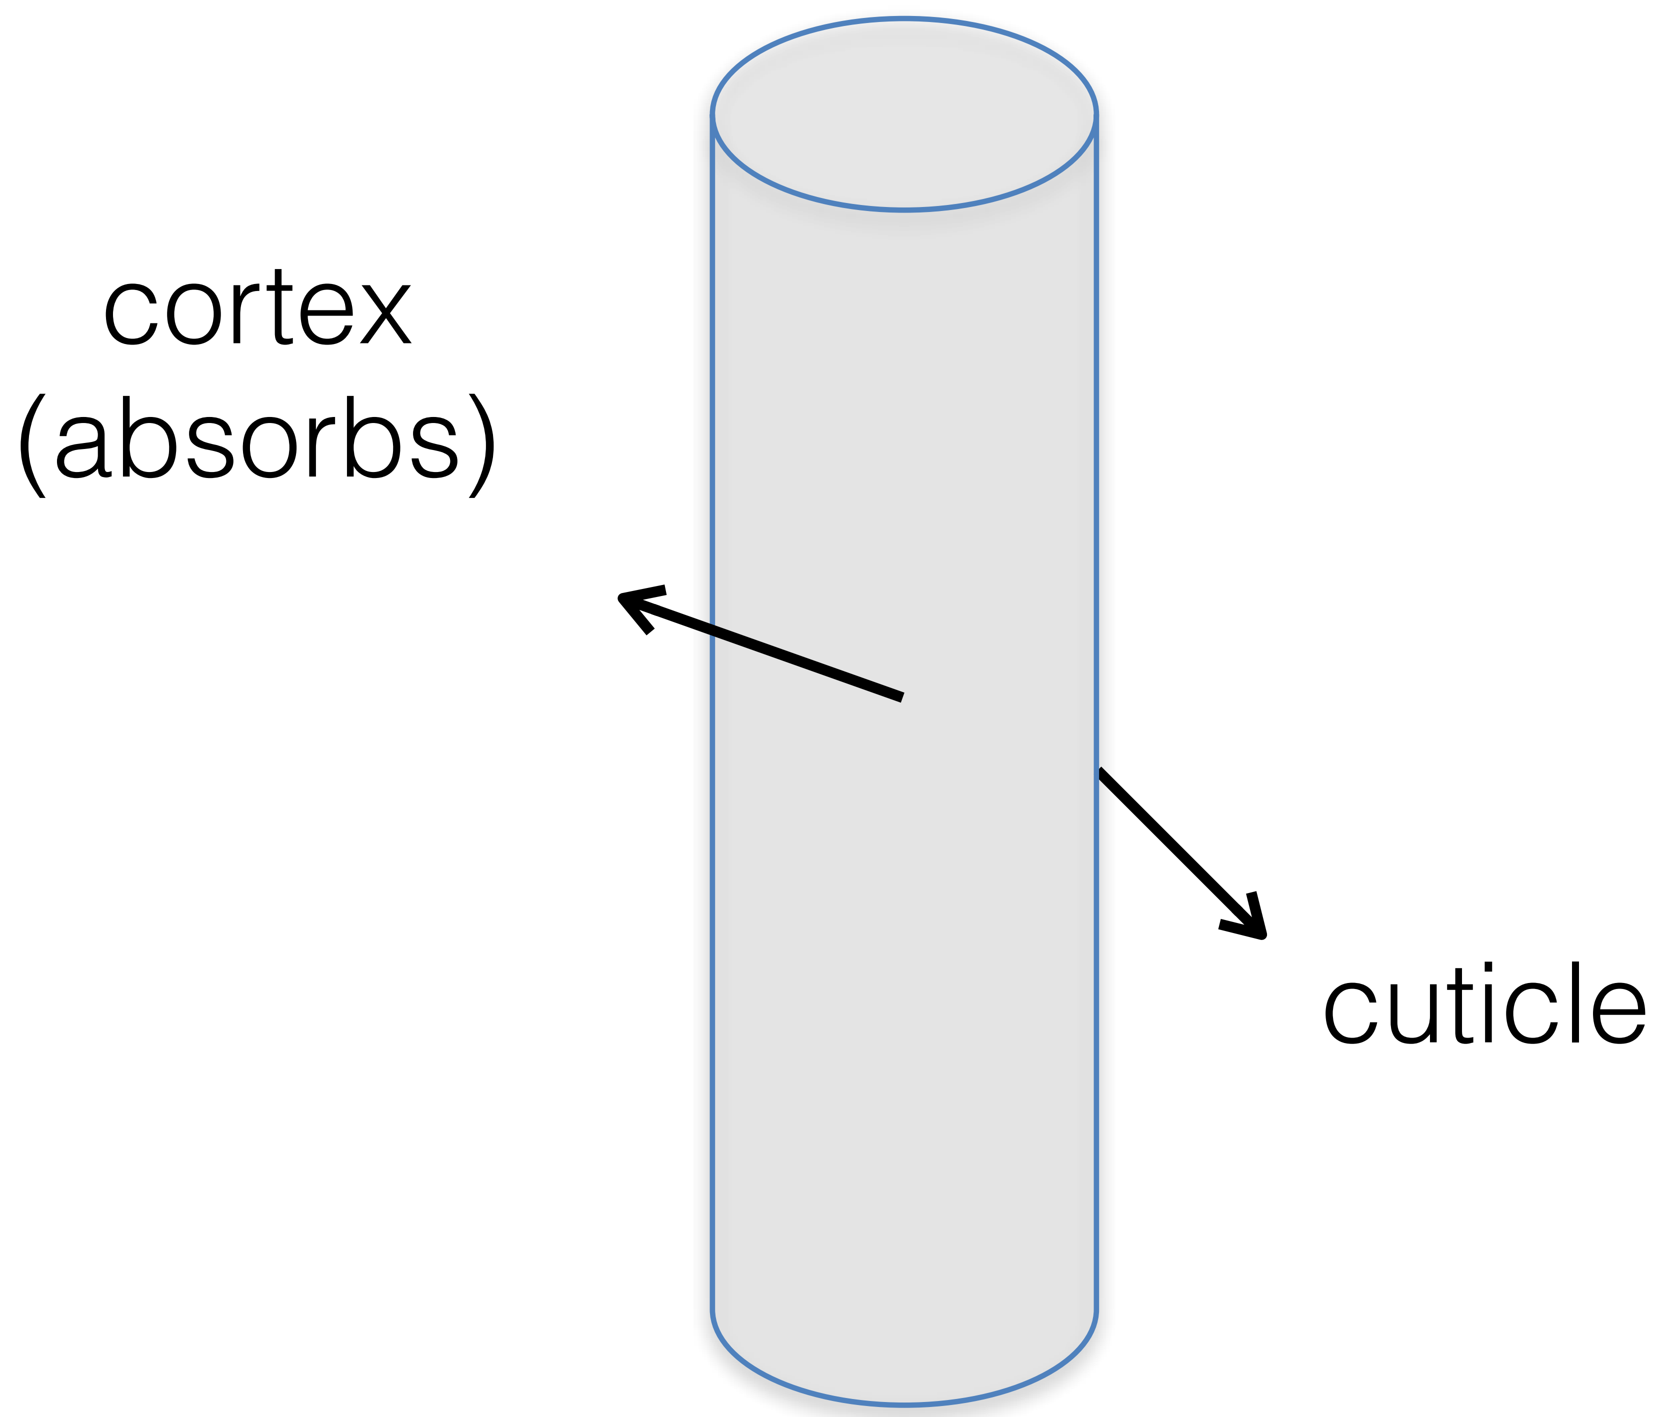

- 3 types of light interactions:
R, TT, TRT
(R: reflection, T: transmission)

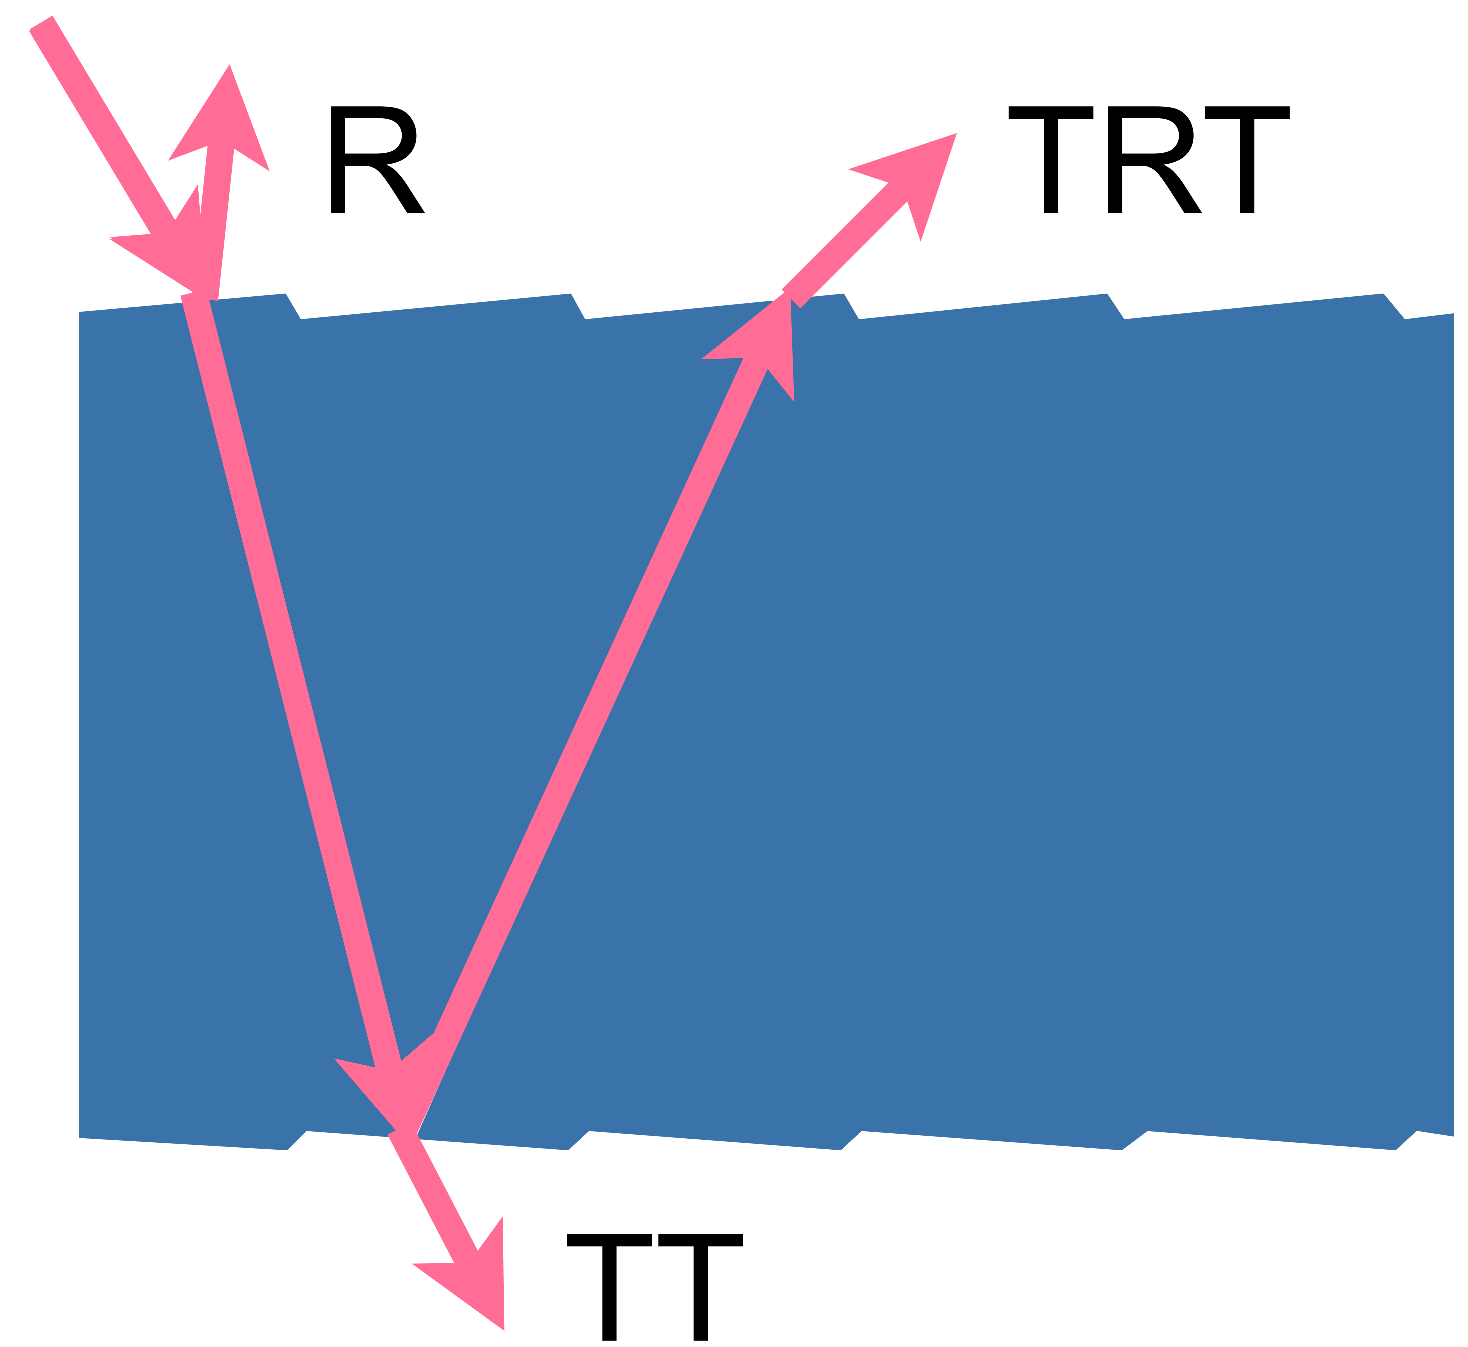

[Marschner et al. 2003]

Marschner model

[Marschner et al. 2003]

[d’Eon et al. 2011]

Hair Appearance Model: Application

[Final Fantasy XV. 2016 Square Enix]

Hair Appearance Model: Application

[Zootopia. 2016 Disney]

Fur Appearance — As Human Hair

- Cannot represent diffusive and saturated appearance

Rendered as human hair
[Marschner et al. 2003]

Rendered as animal fur
[Yan et al. 2015]

Human Hair vs Animal Fur

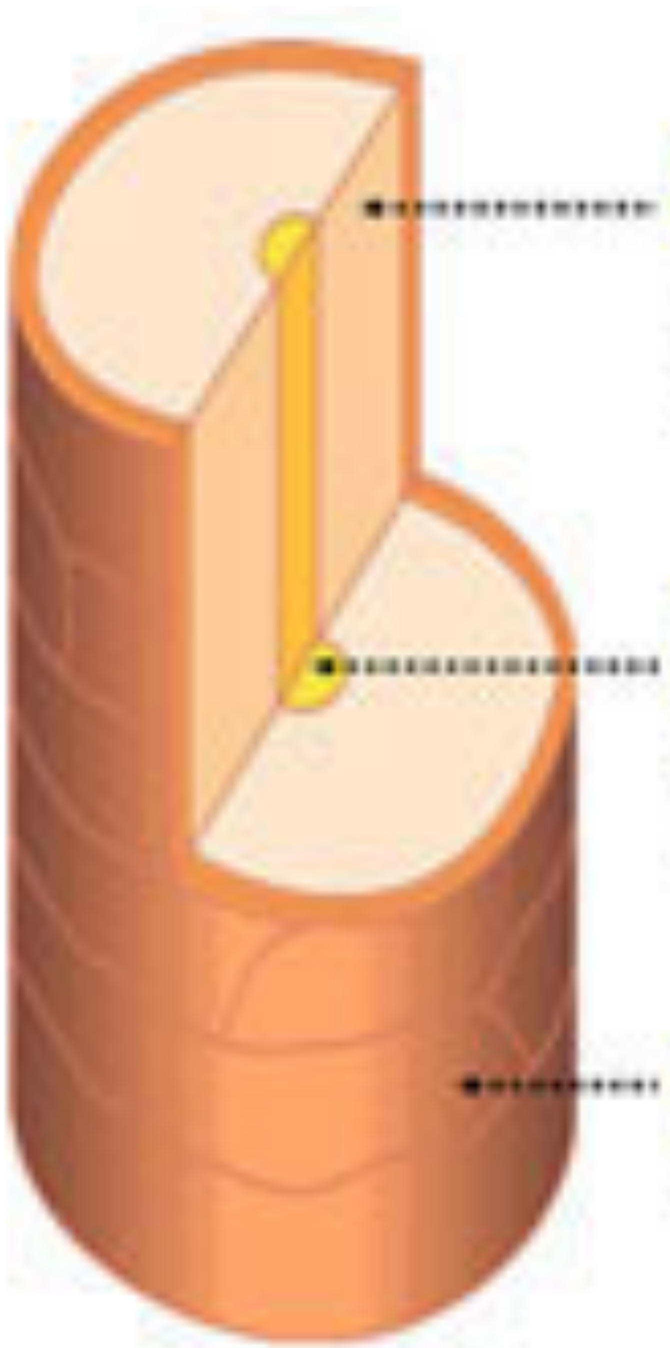

Cortex
- Contains pigments
- Absorbs light

Medulla
- Complex structure
- Scatters light

Cuticle
- Covered with scales

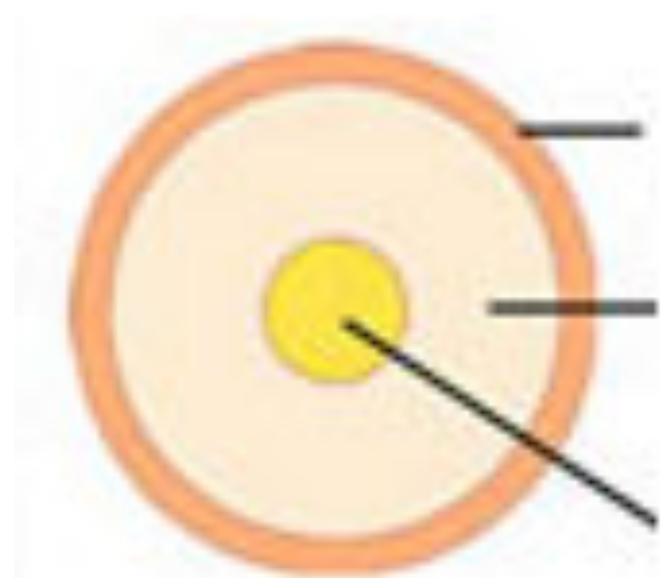

Common for hair/fur fibers

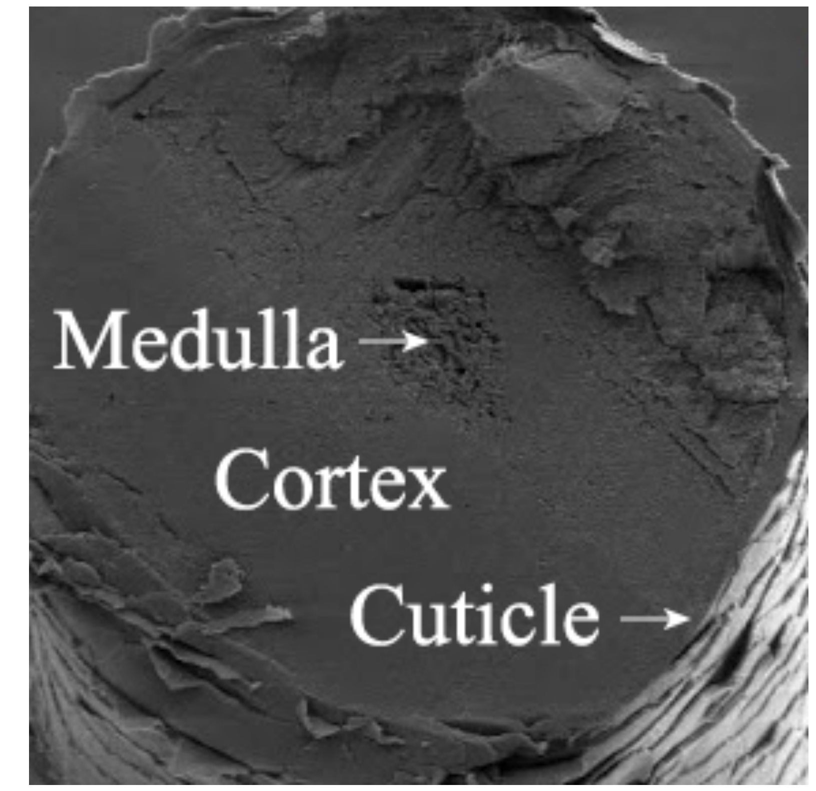

Human

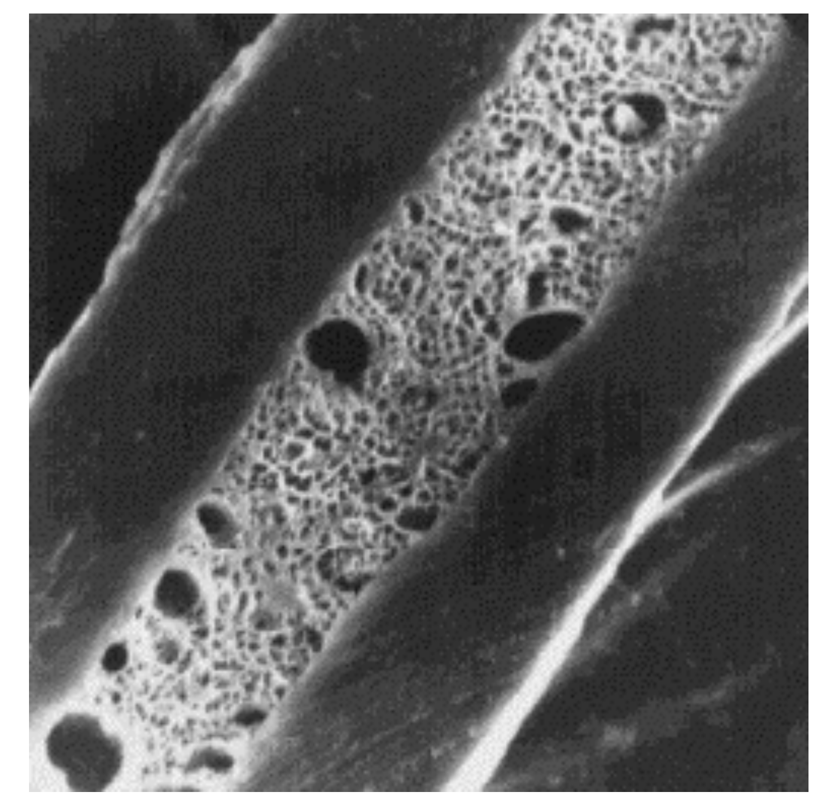

Cougar

Difference between hair/fur fibers

Importance of Medulla

Increasing medulla size

Importance of Medulla

Without medulla

With medulla (**15%**)

Double Cylinder Model

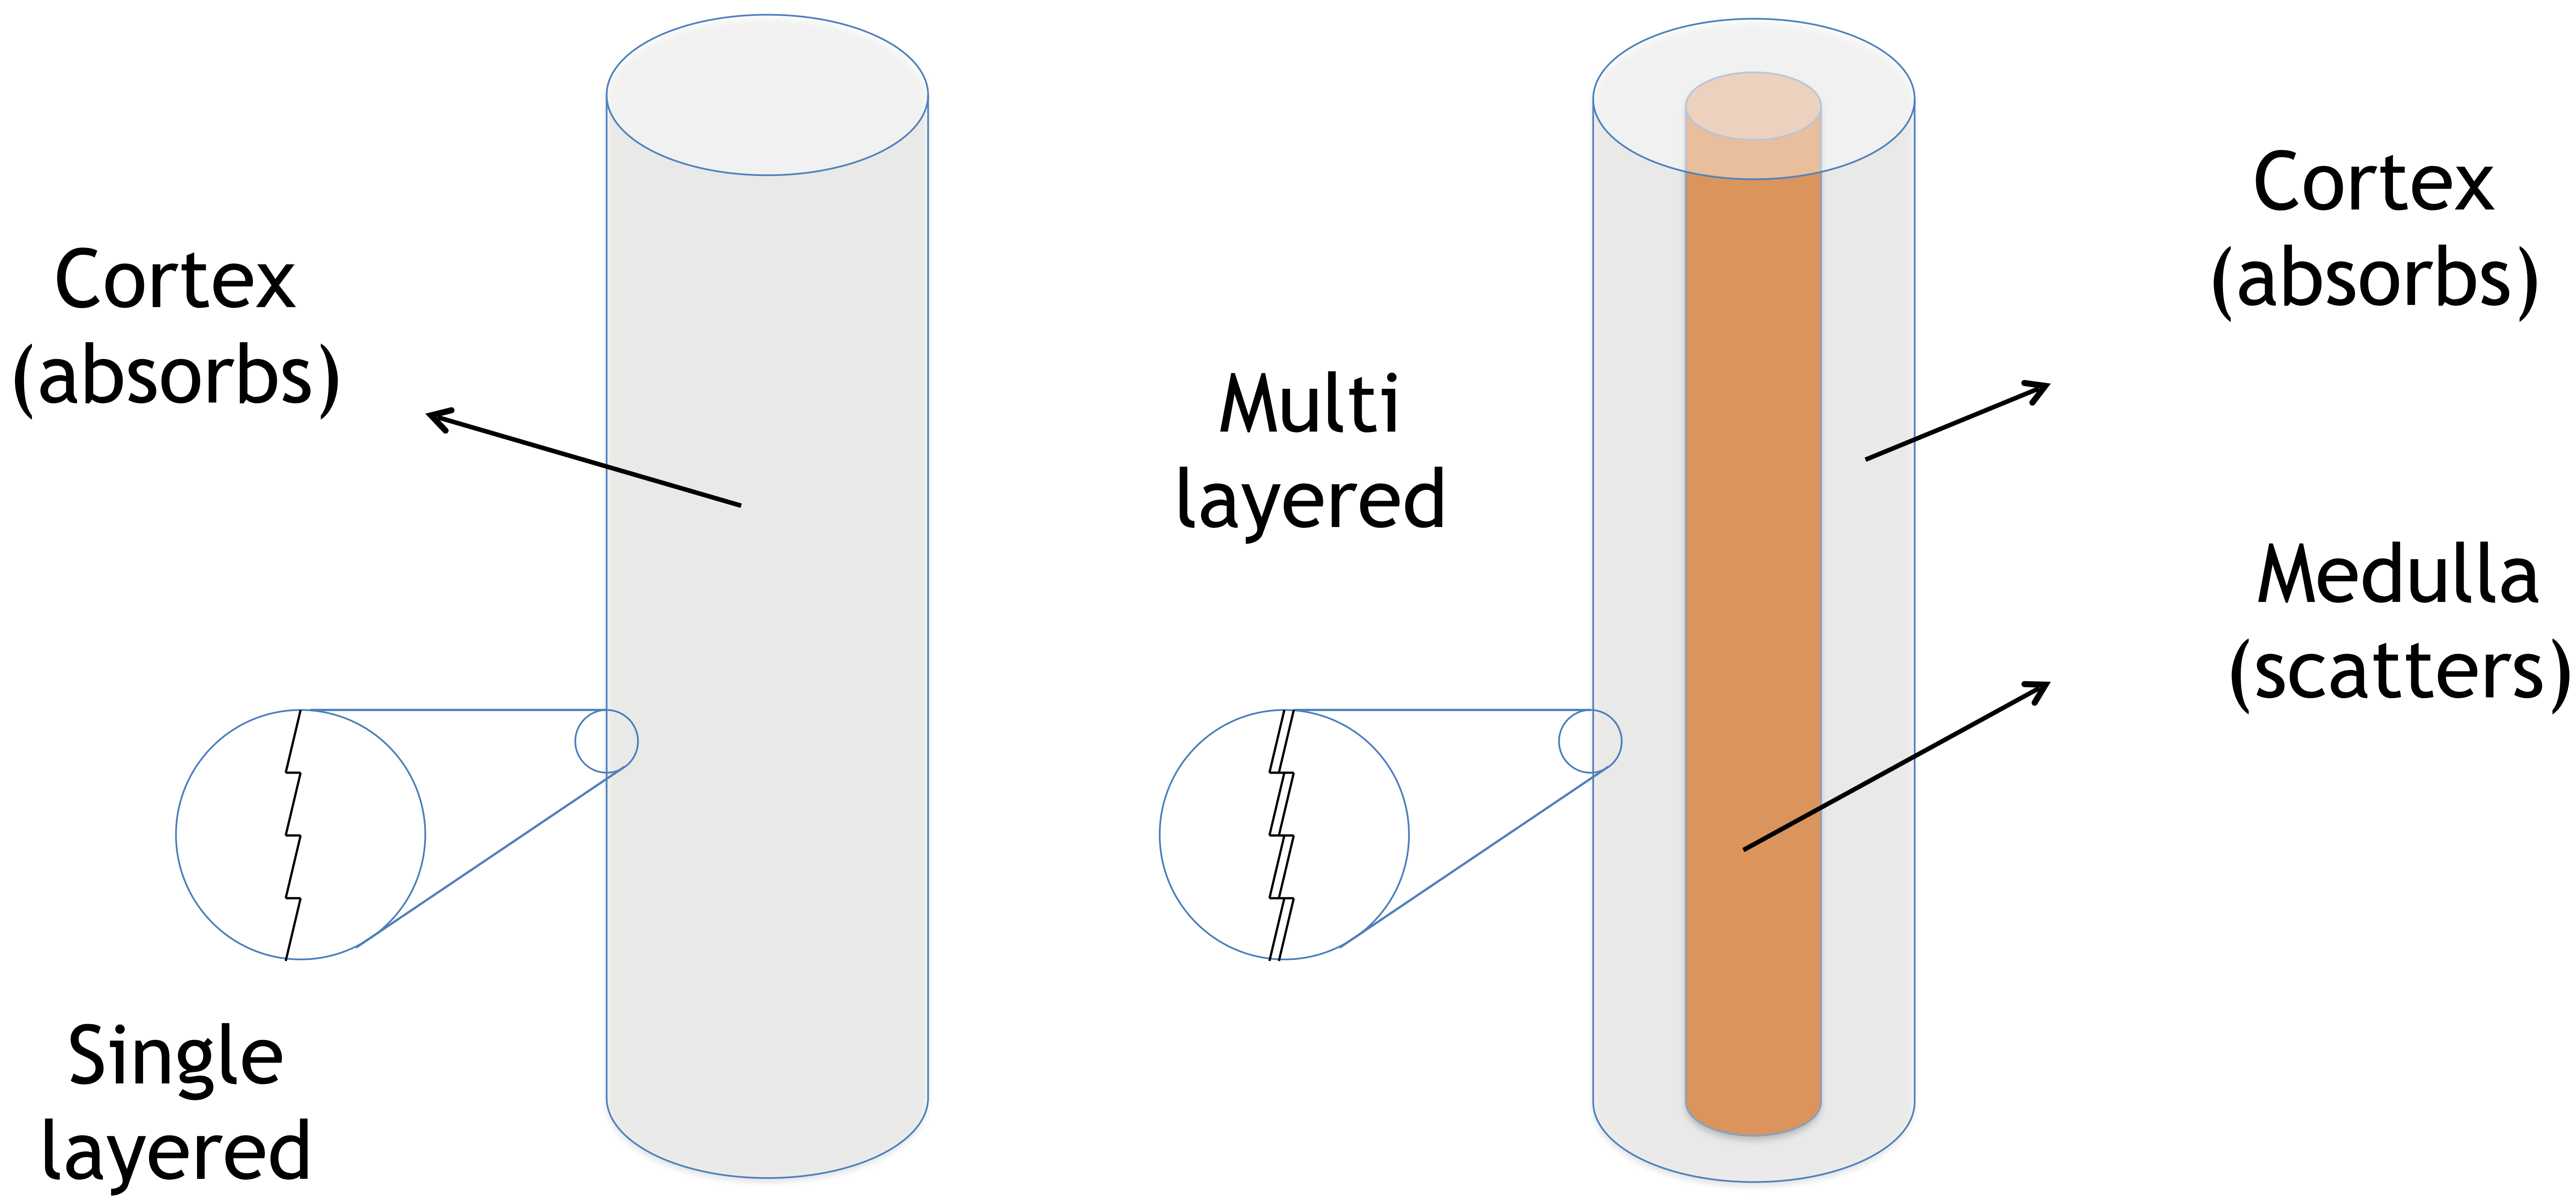

Marschner Model

Double Cylinder Model

[Yan et al. 2015, 2017]

Double Cylinder Model — Lobes

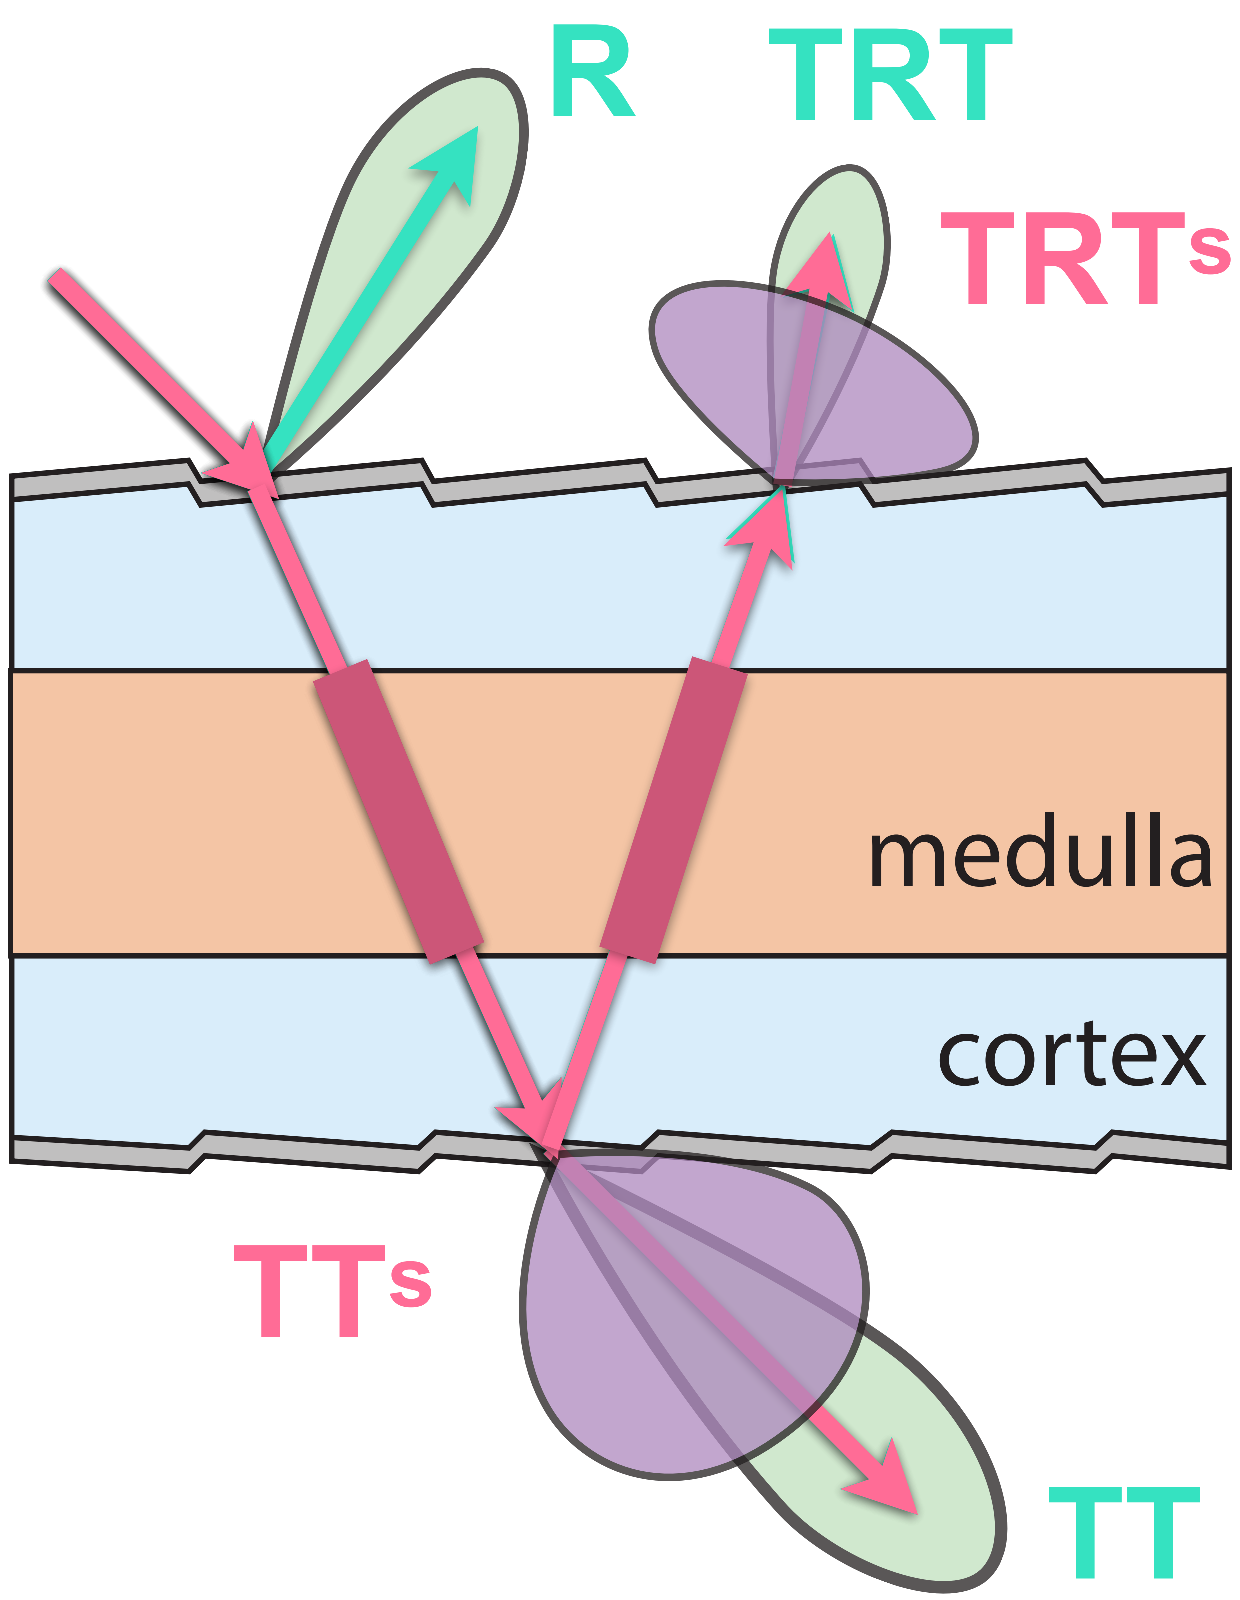

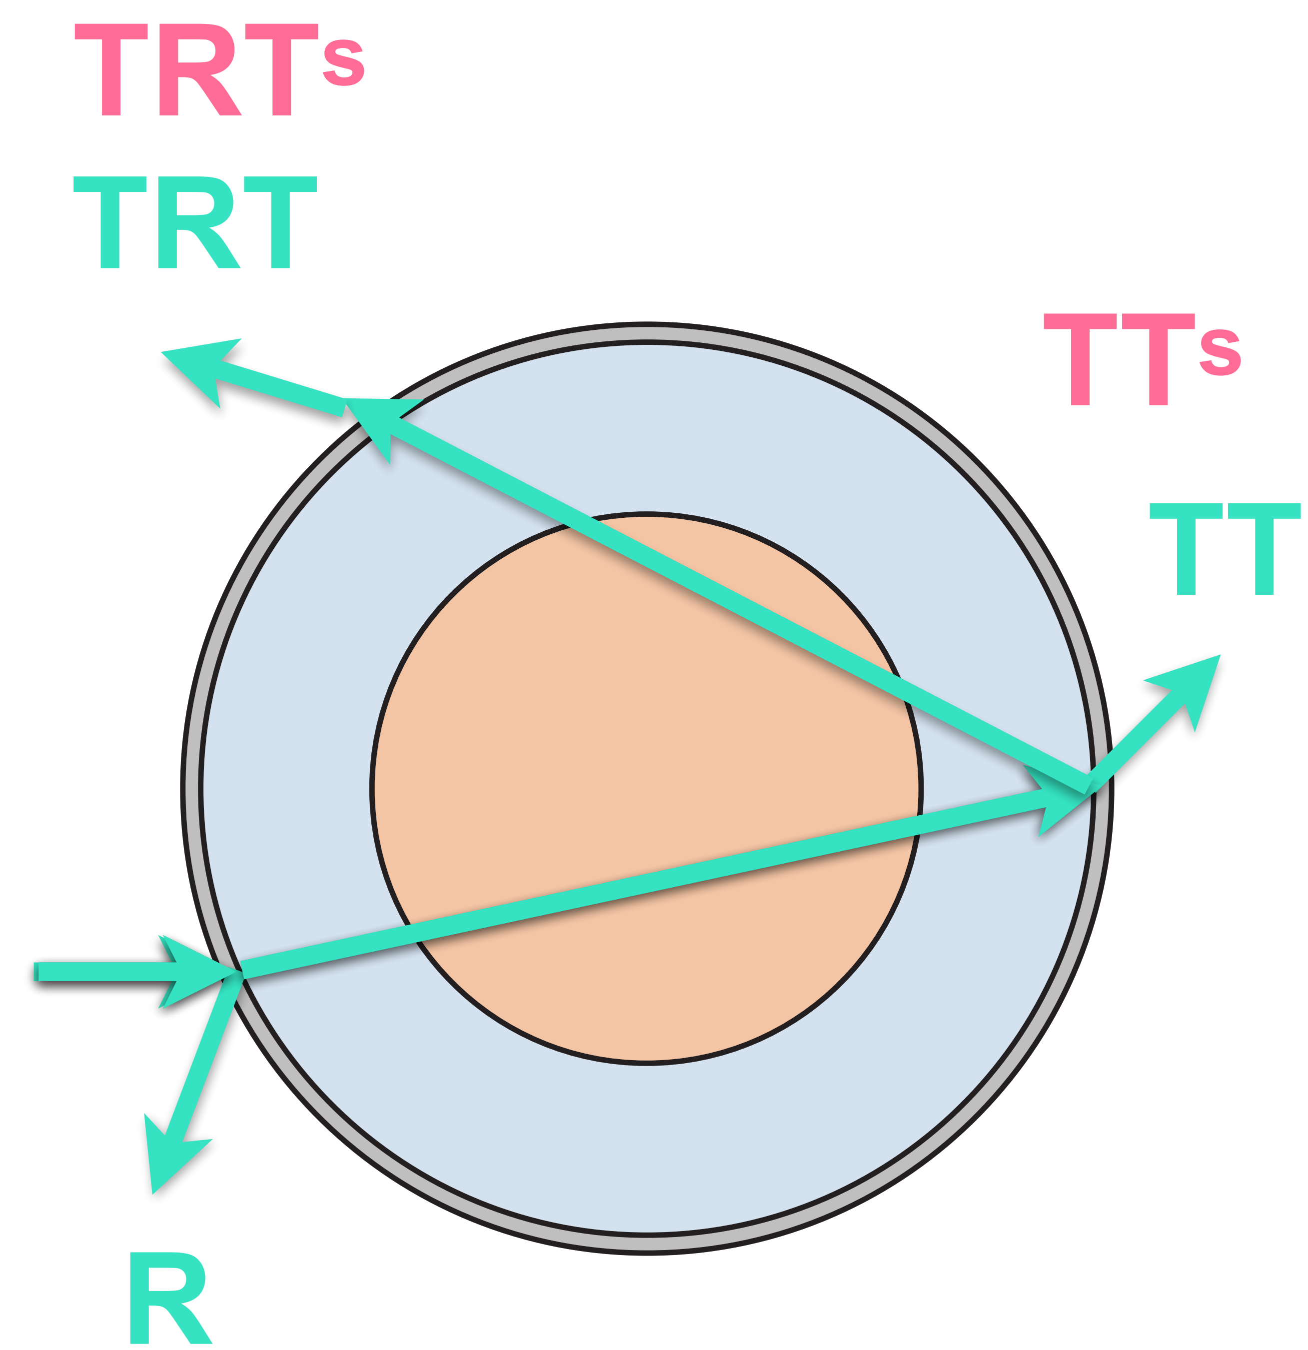

Double Cylinder Model — Lobes

All

=

R TT TRT TTs TRTs

600,000 fur fibers
1024 samples / pixel
36.9 min / frame

Hamster

Double Cylinder Model: Application

[War for the Planet of the Apes. 2017 movie] (2018 Oscar Nominee for Best Visual Effects)

Double Cylinder Model: Application

[The Lion King (HD). 2017 movie] (2019 Oscar Nominee for Best Visual Effects)

Granular Material

- What is granular material?

[Meng et al. 2015]

Granular Material

- Can we avoid explicit modeling of all granules?
  - Yes with procedural definition.

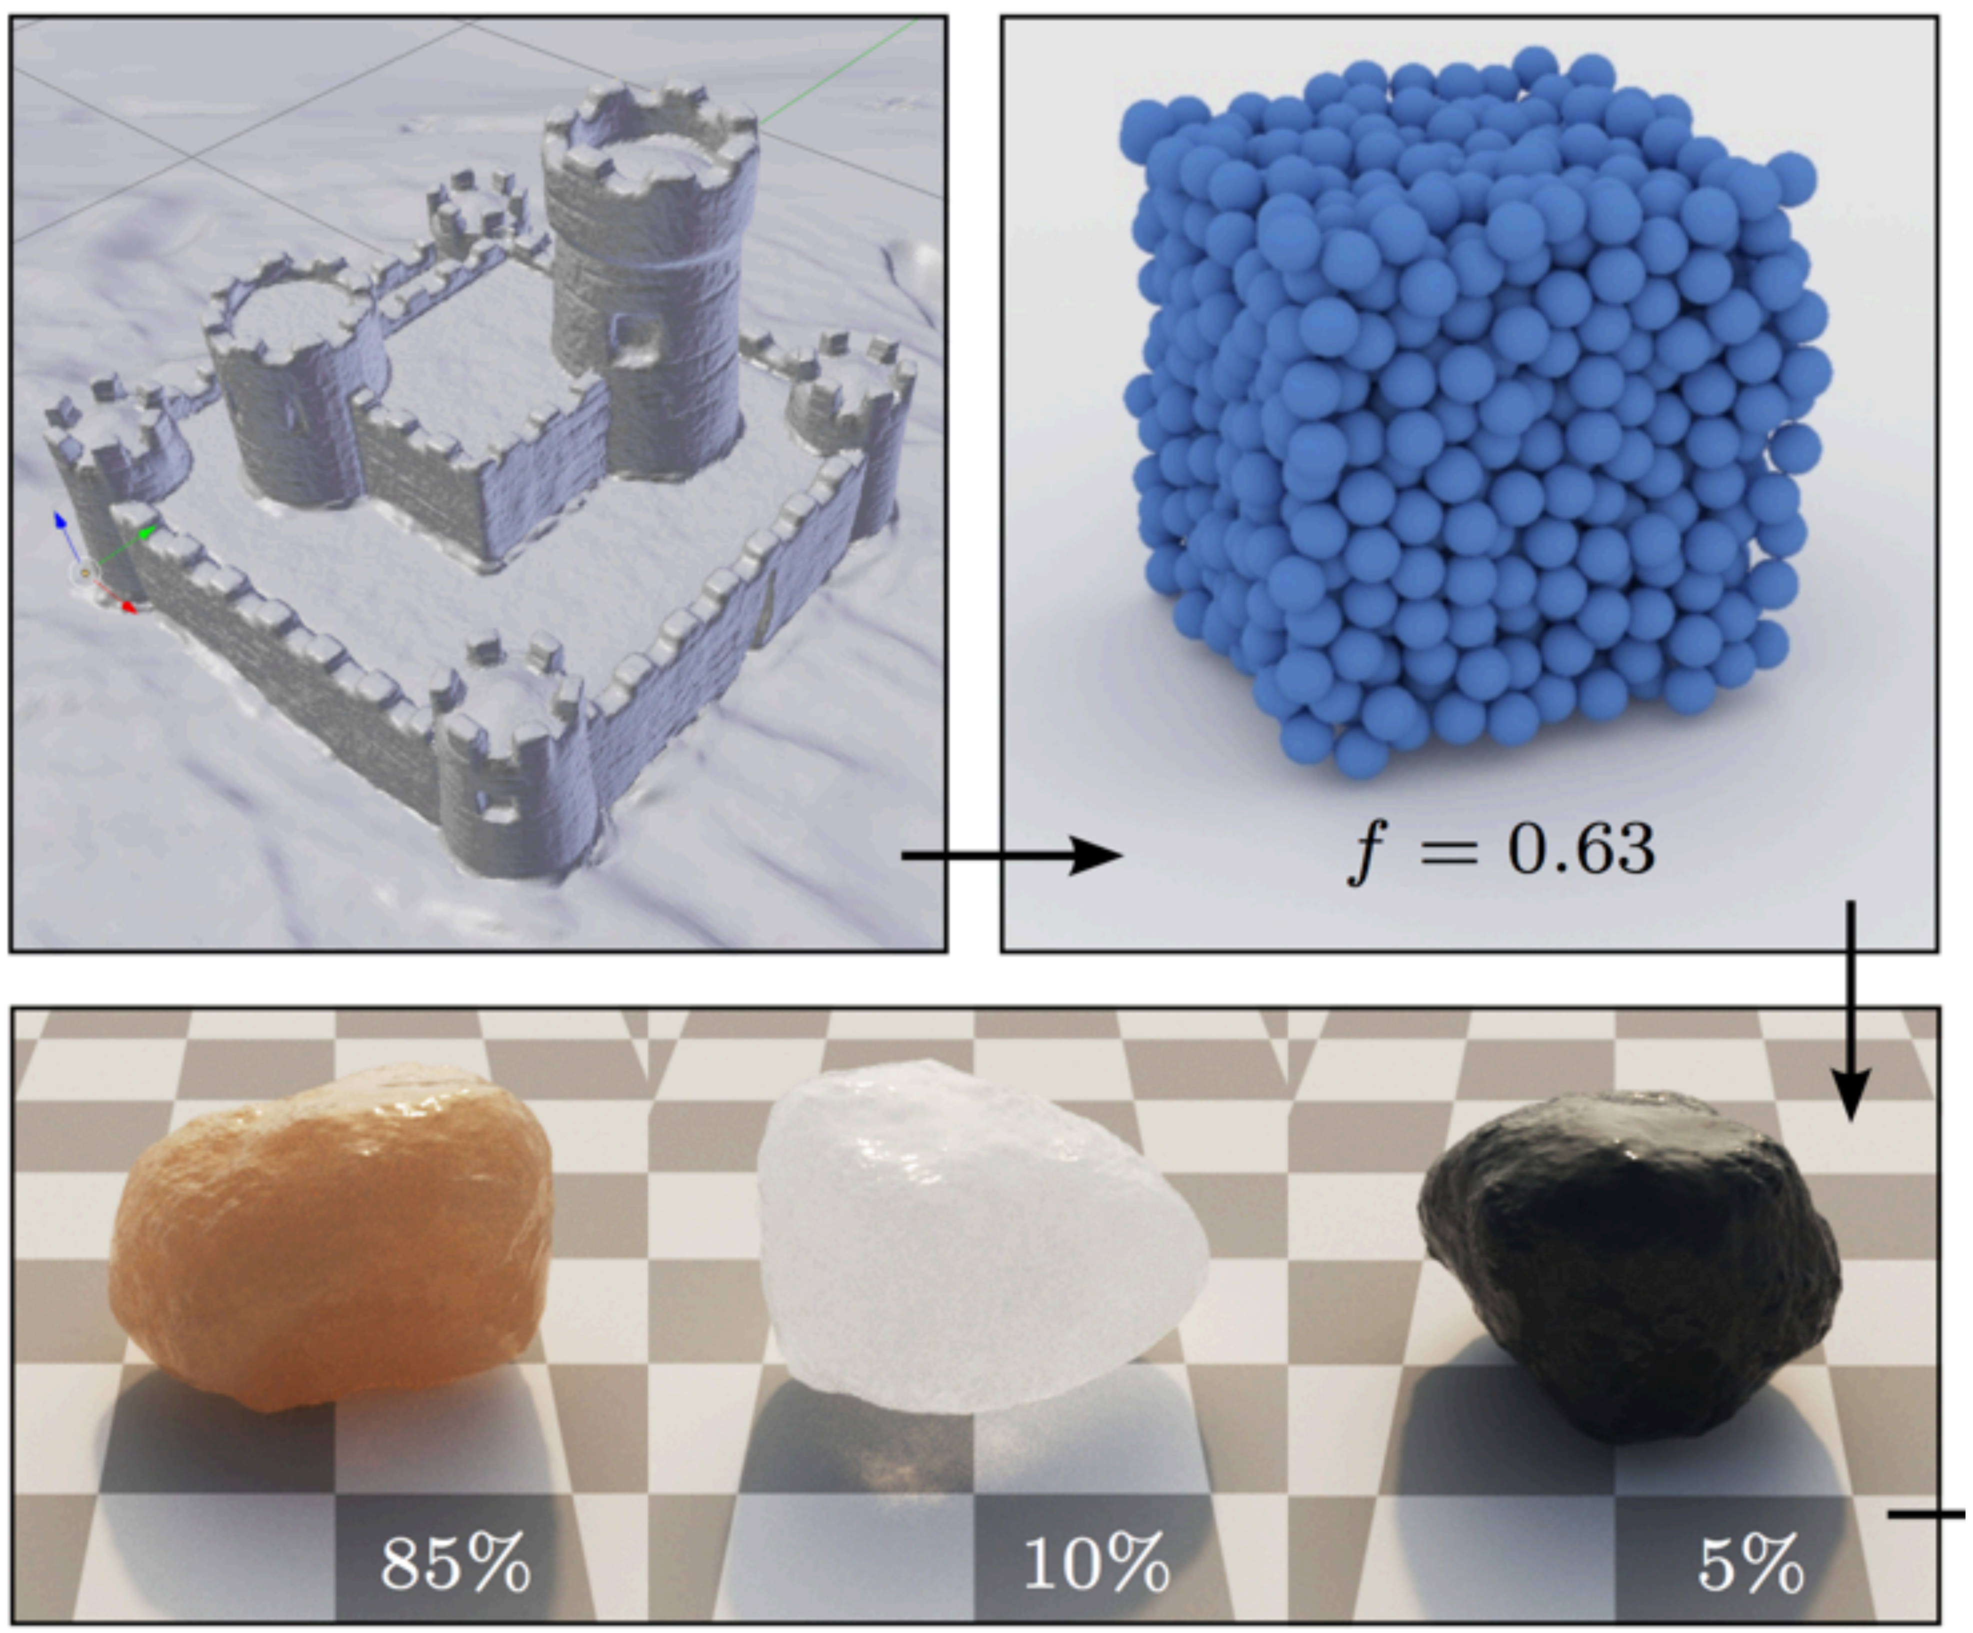

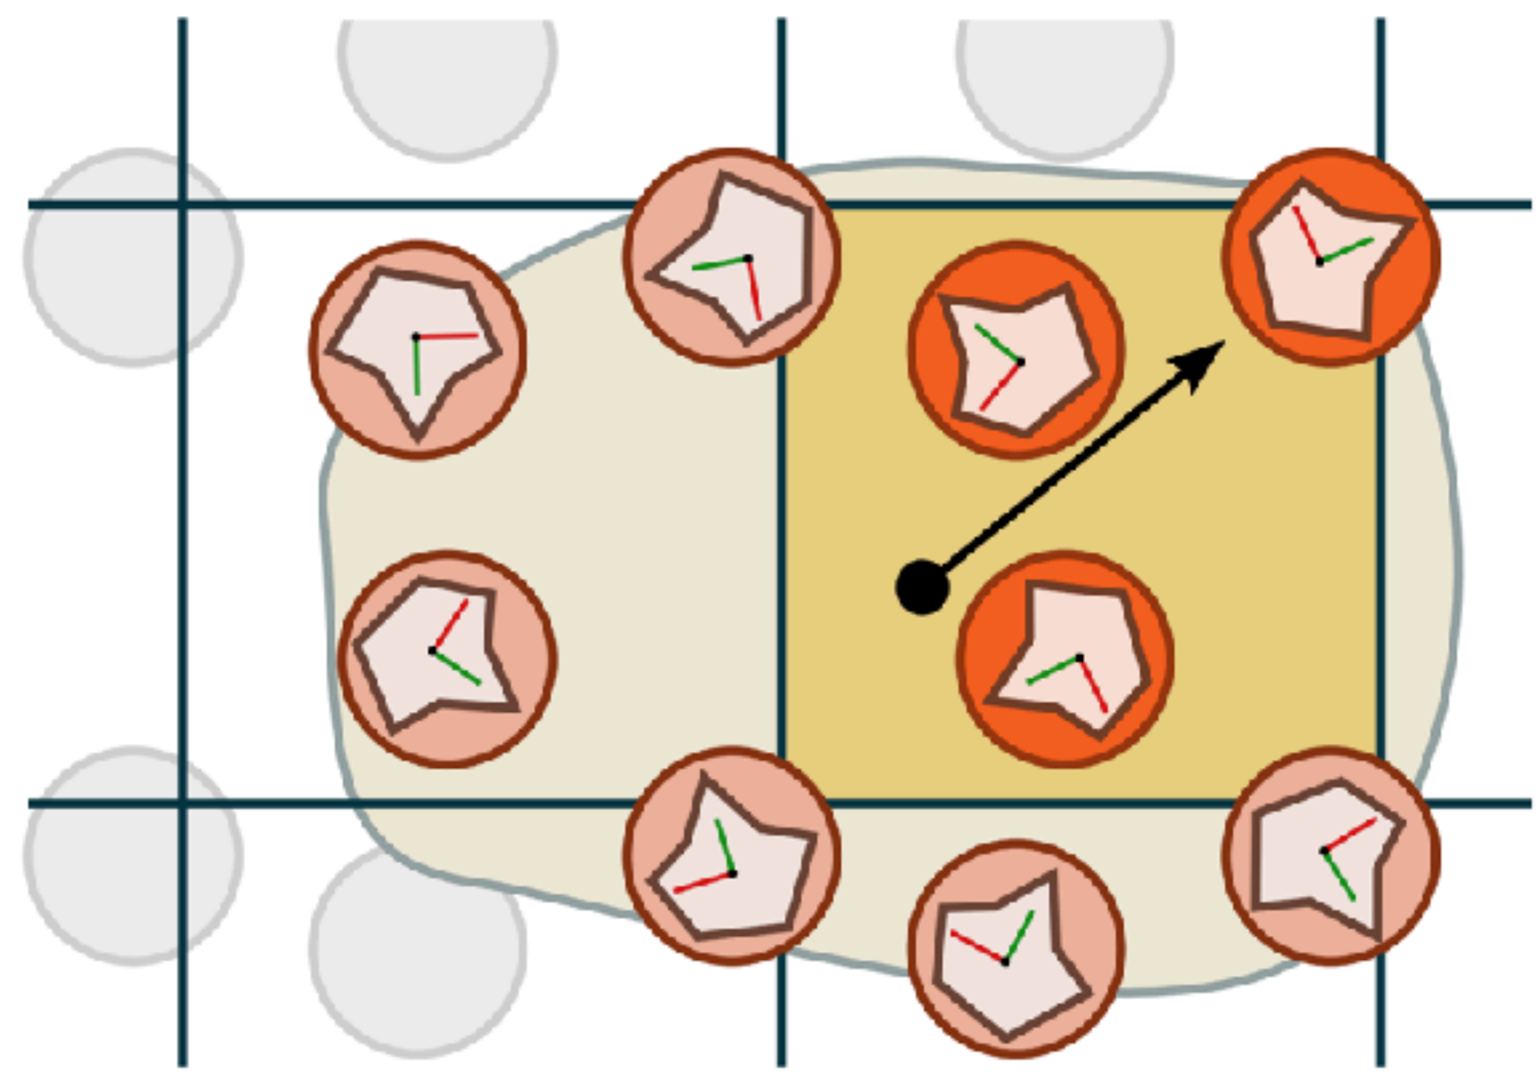

[Meng et al. 2015]

Granular Material

[Meng et al. 2015]

Granular Material: Application

[Piper. 2016 Pixar]

Surface Models

Translucent Material: Jade

Translucent Material: Jellyfish

Subsurface Scattering

Visual characteristics of many surfaces caused by light exiting at different points than it enters

- Violates a fundamental assumption of the BRDF

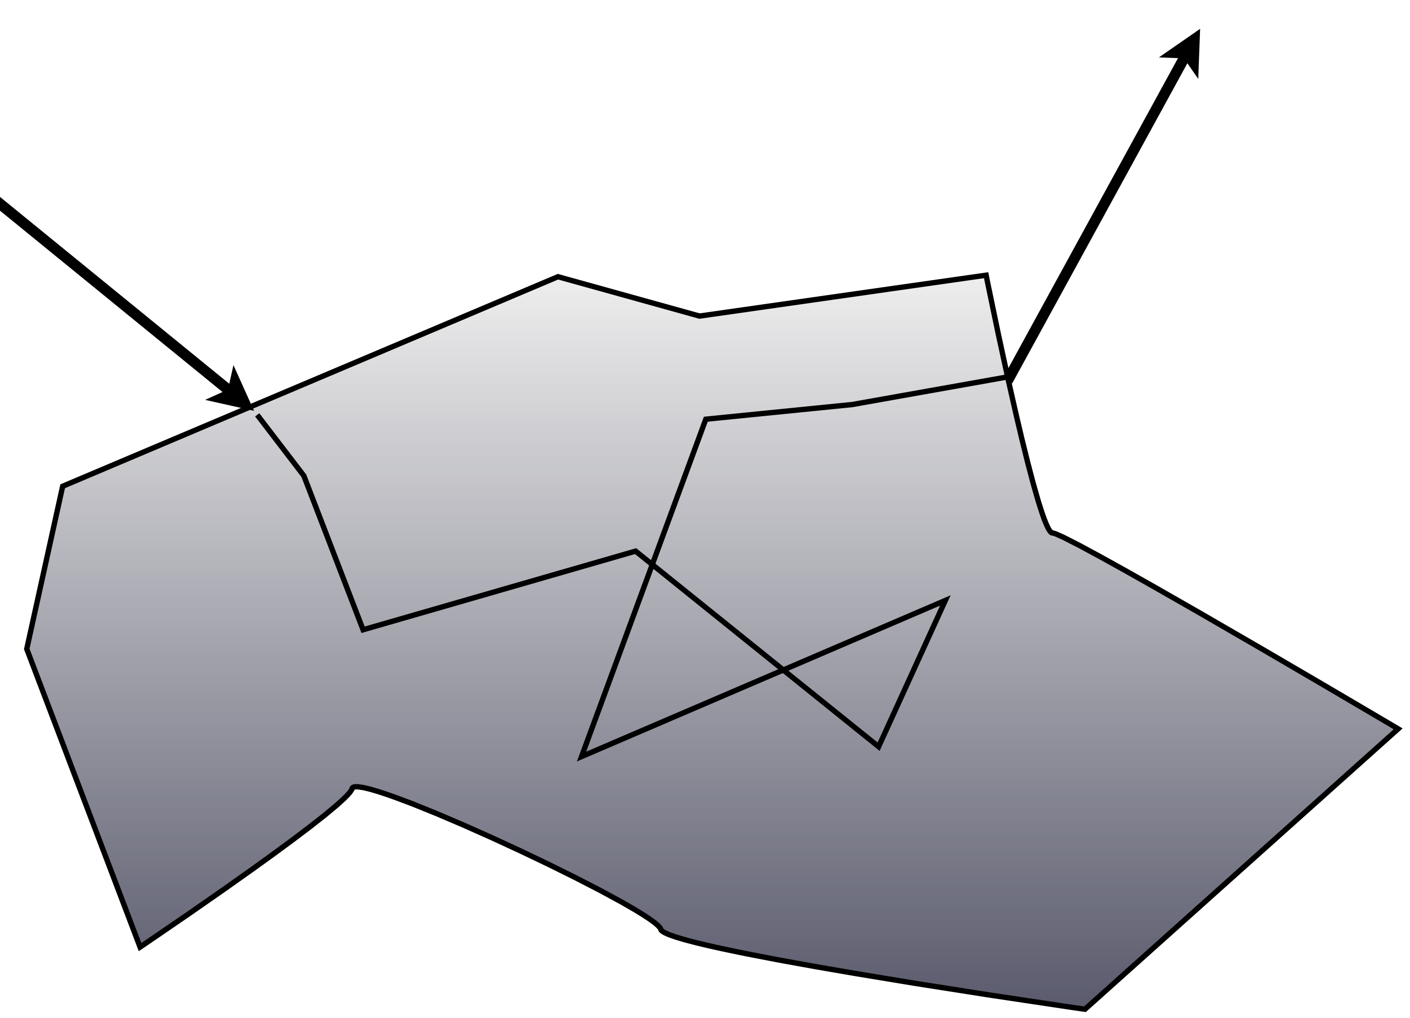

[Jensen et al 2001]

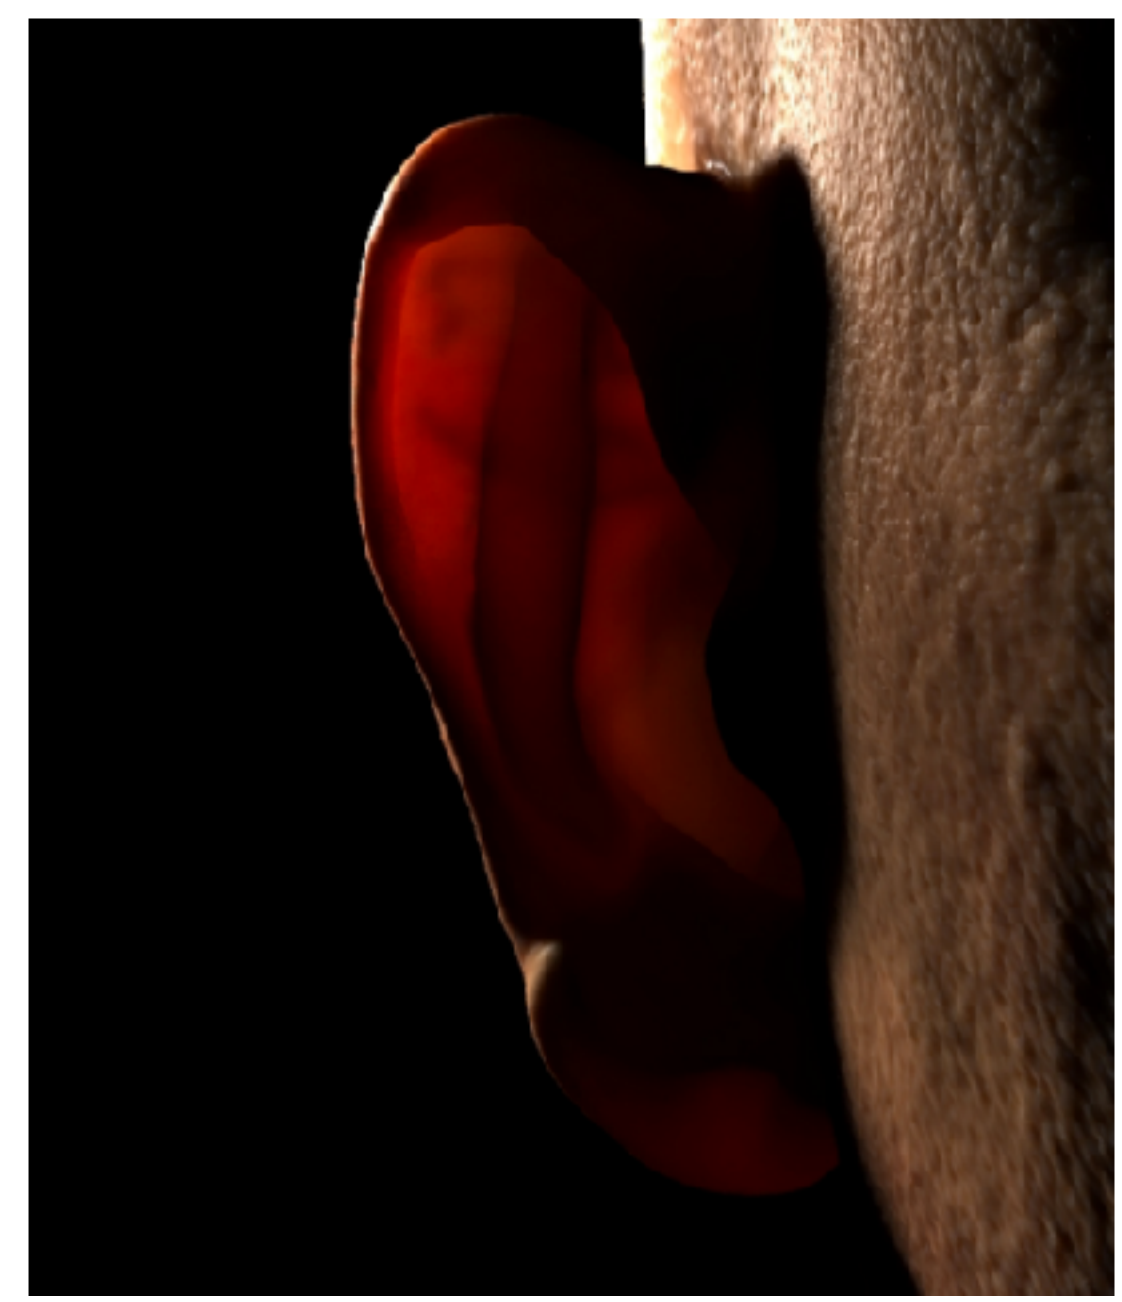

[Donner et al 2008]

Scattering Functions

- BSSRDF: generalization of BRDF; exitant radiance at one point due to incident differential irradiance at another point:

$$S(x_i, \omega_i, x_o, \omega_o)$$

- Generalization of rendering equation: integrating over all points on the surface and all directions (!)

$$L(x_o, \omega_o) = \int_A \int_{H^2} S(x_i, \omega_i, x_o, \omega_o)\, L_i(x_i, \omega_i) \cos\theta_i \,\mathrm{d}\omega_i \,\mathrm{d}A$$

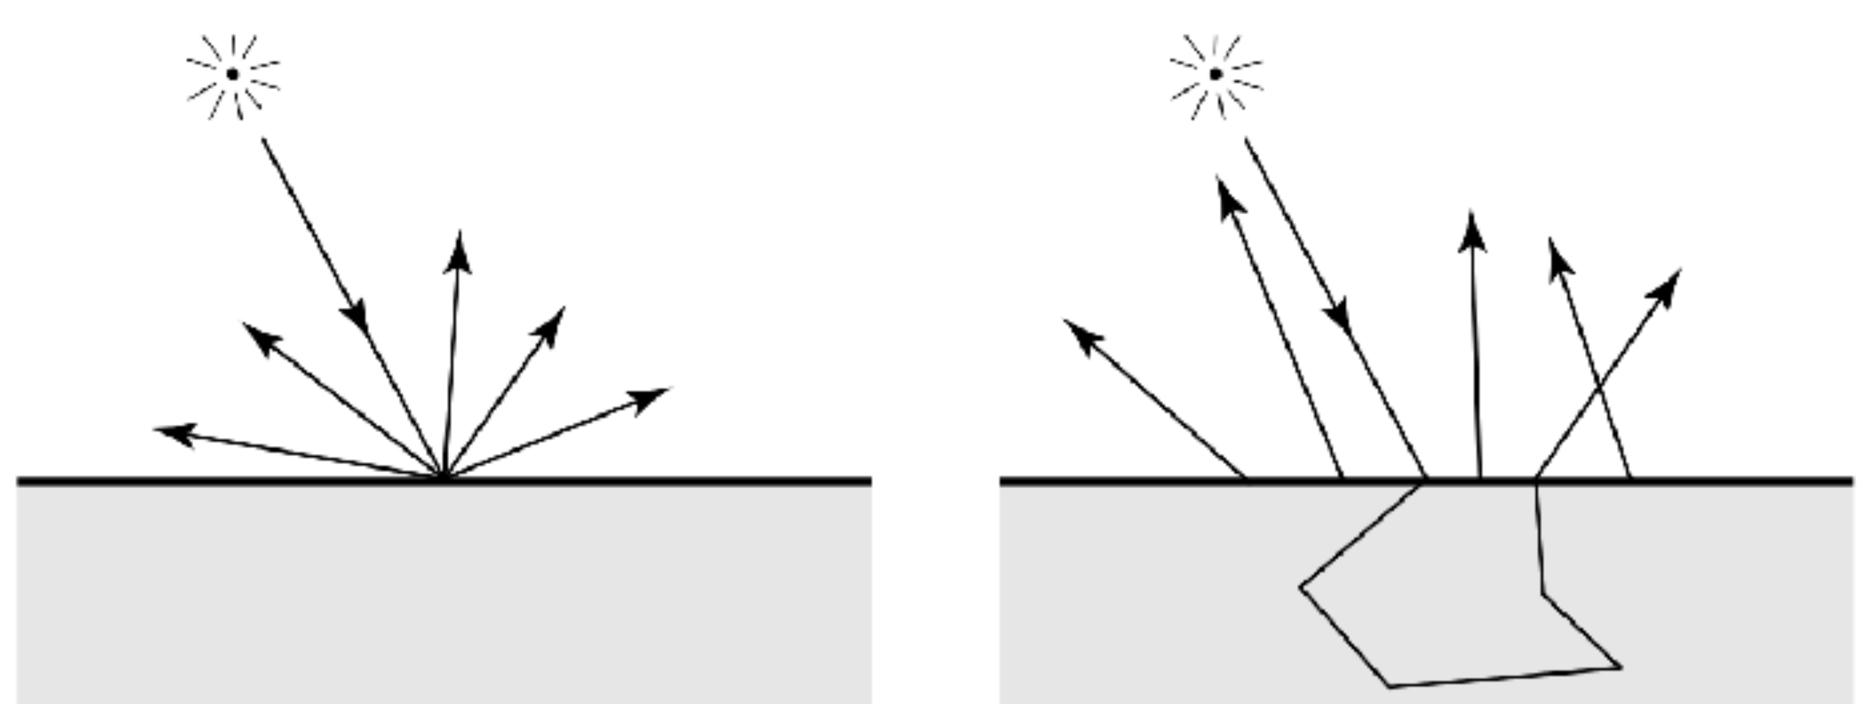

BRDF

BSSRDF

[Jensen et al. 2001]

Dipole Approximation [Jensen et al. 2001]

- Approximate light diffusion by introducing two point sources.

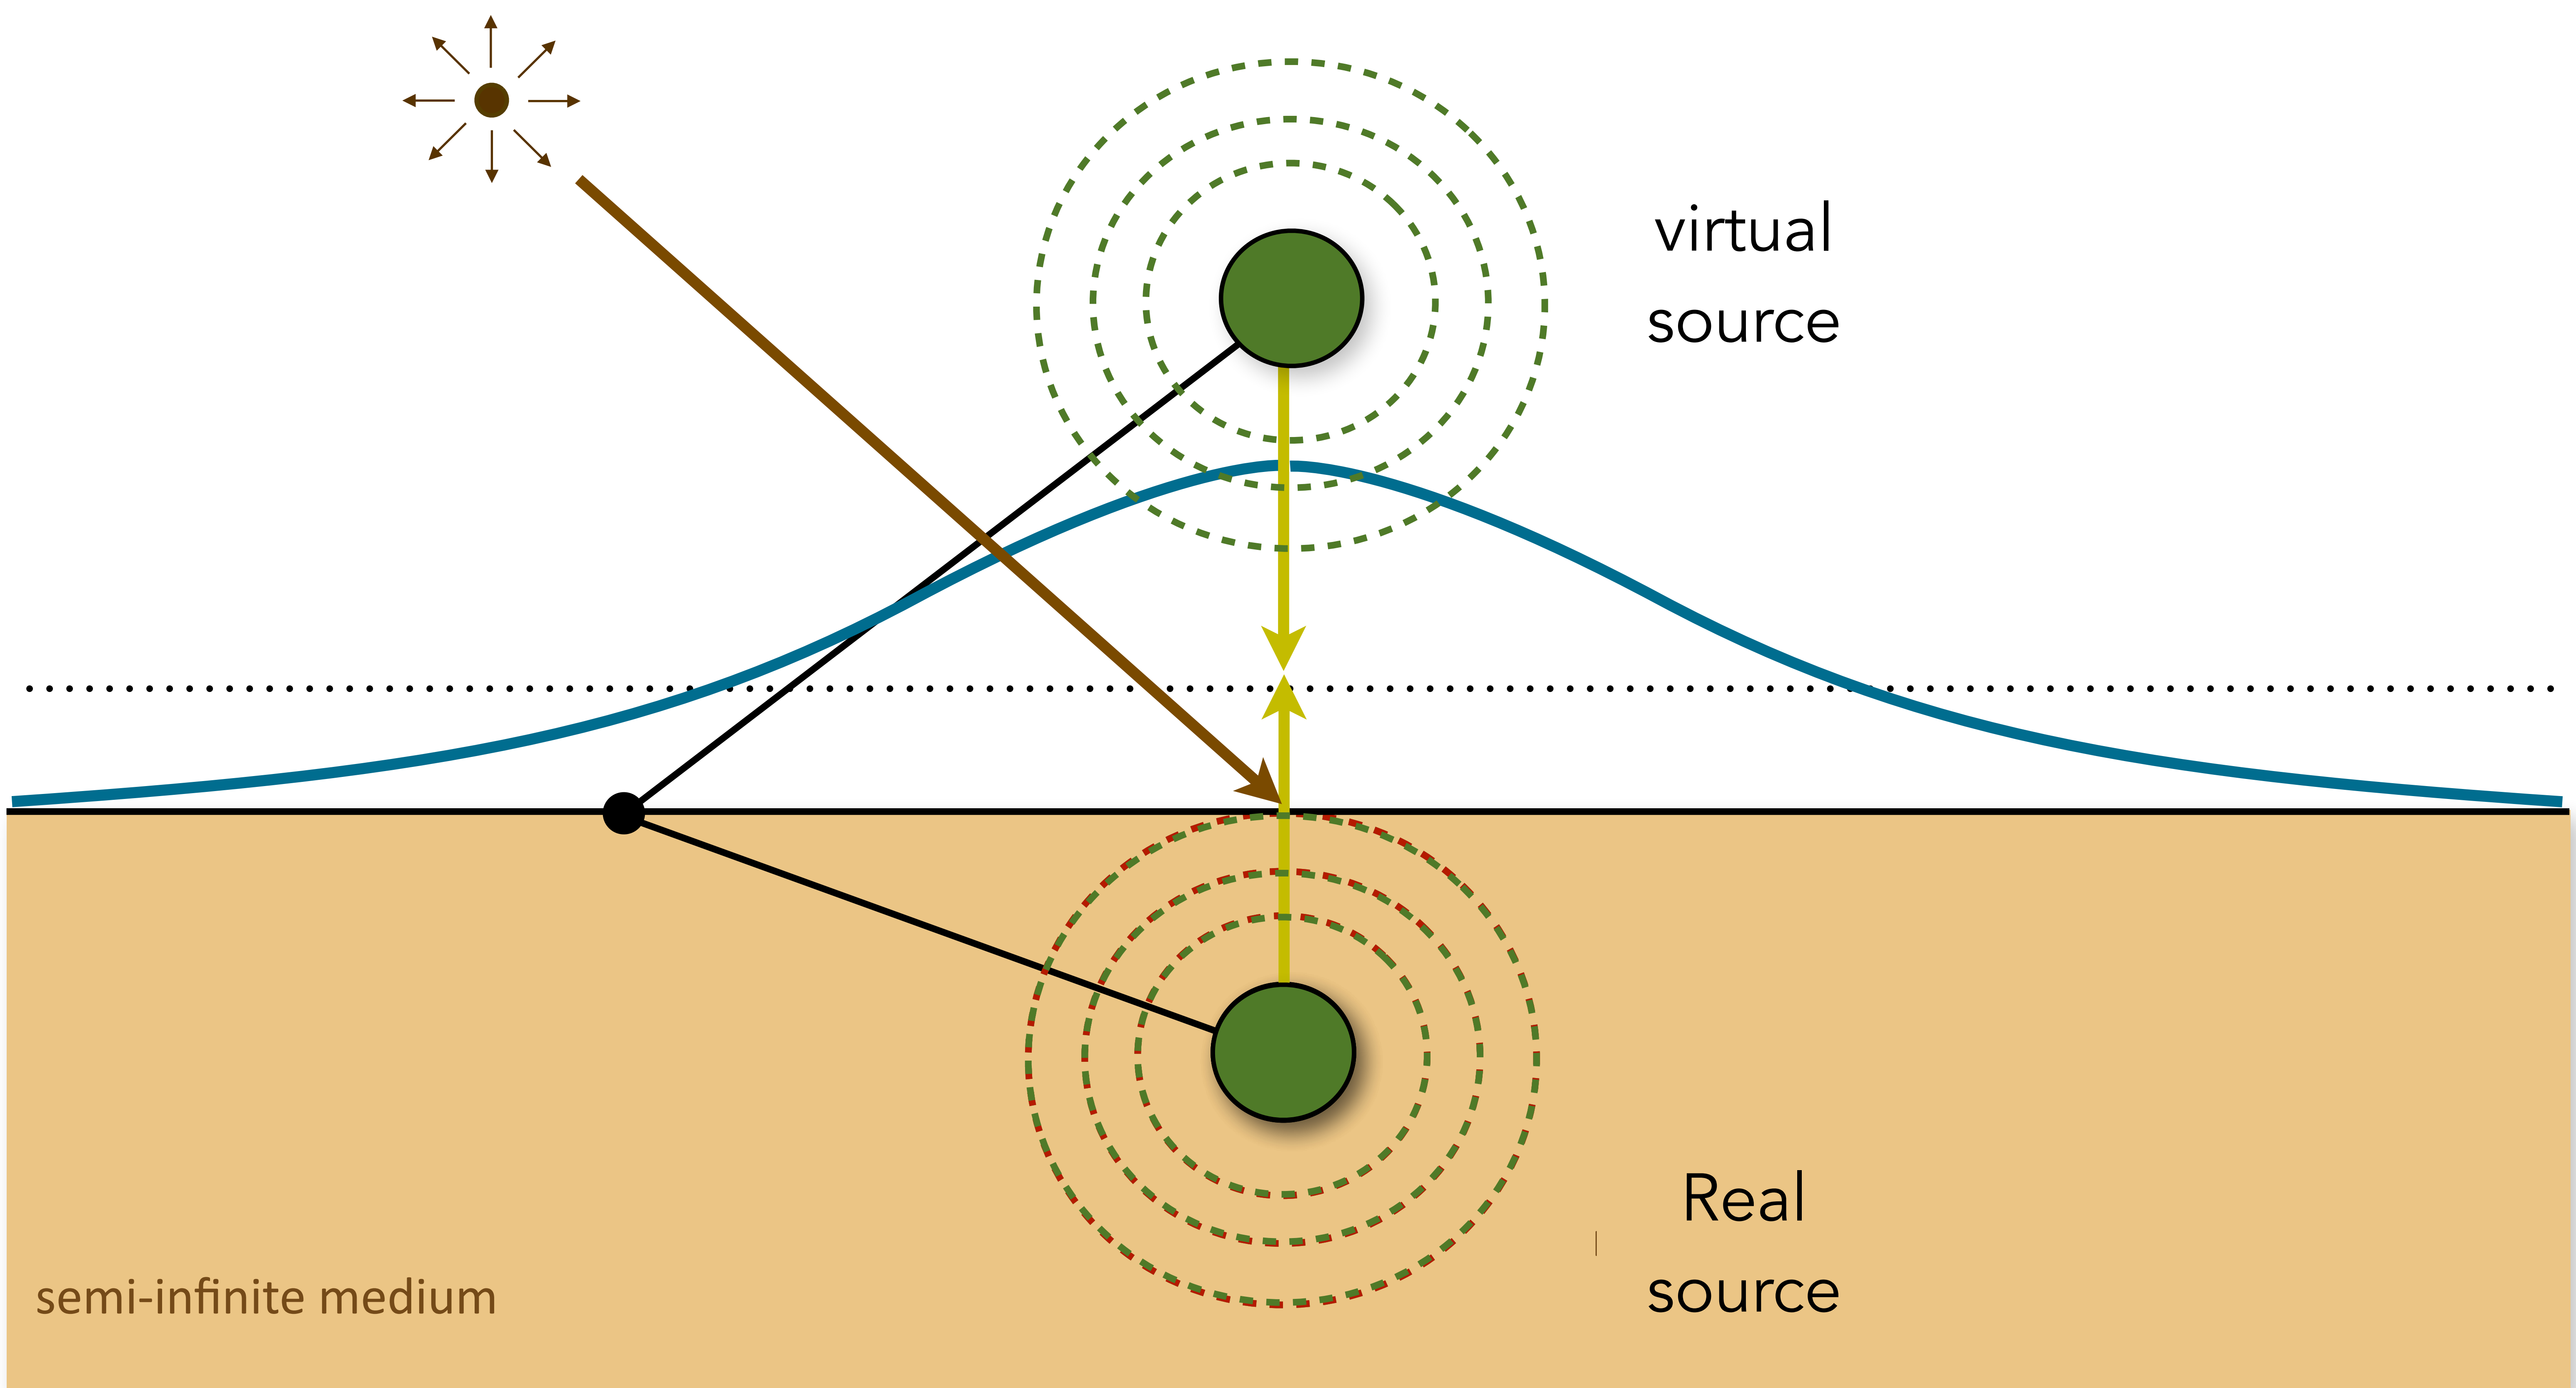

[image from Habel et al. 2013]

BRDF

[Jensen et al. 2001]

BSSRDF

[Jensen et al. 2001]

BRDF vs BSSRDF

BRDF

BSSRDF

[Jensen et al. 2001]

BSSRDF: Demo

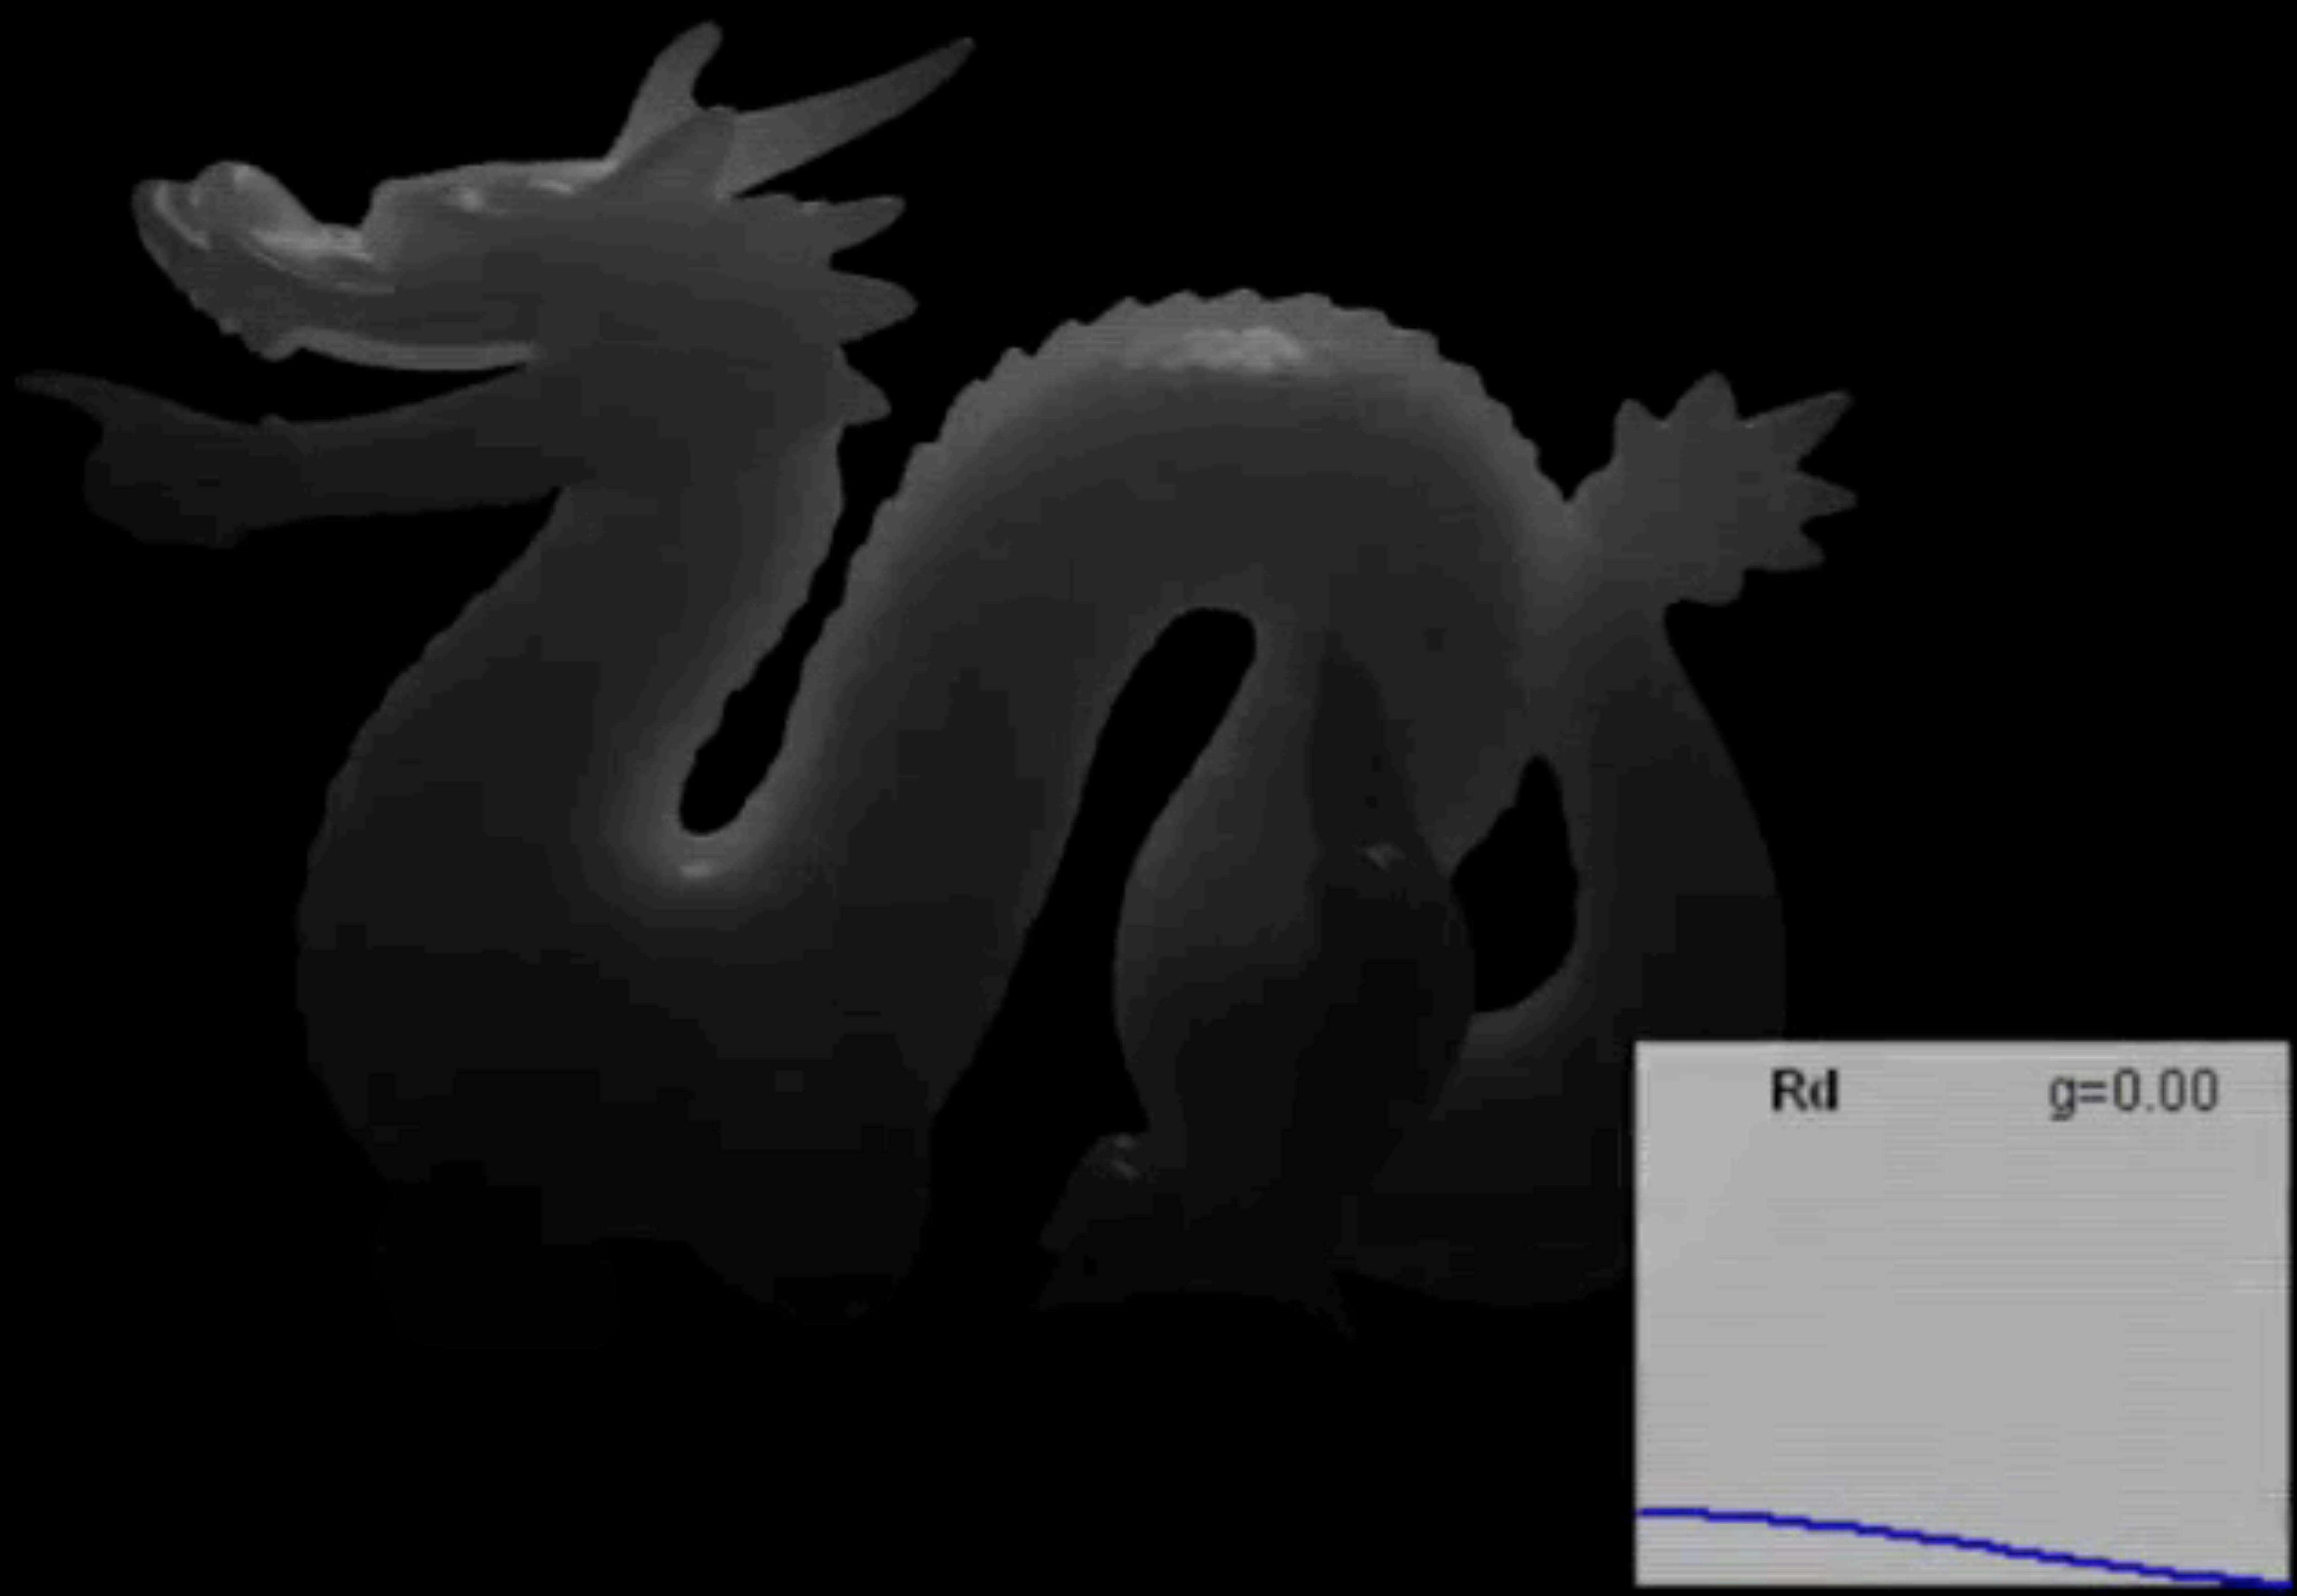

BSSRDF: Application

[Artist: Teruyuki and Yuka]

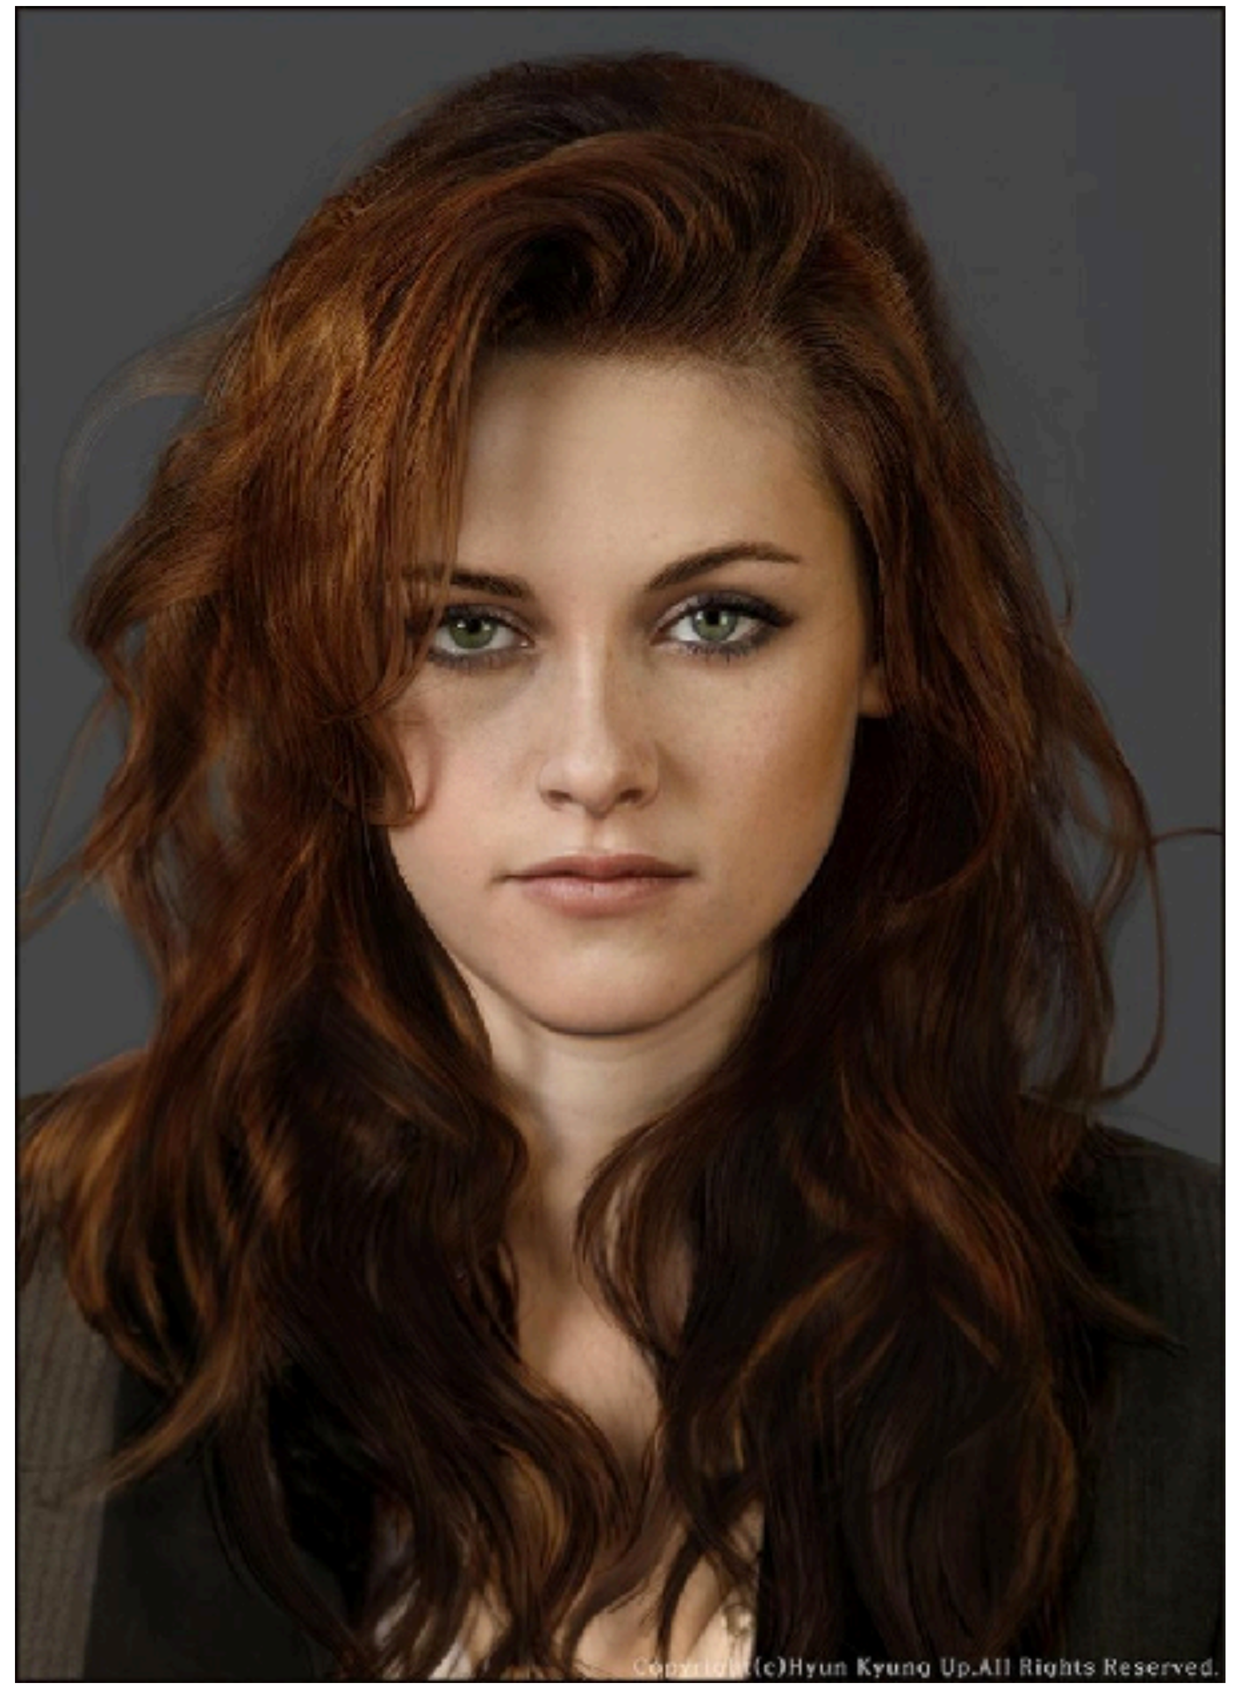

[Artist: Hyun Kyung]

[Artist: Dan Roarty]

https://cgelves.com/10-most-realistic-human-3d-models-that-will-wow-you/

Cloth

- A collection of twisted fibers!
- Two levels of twist

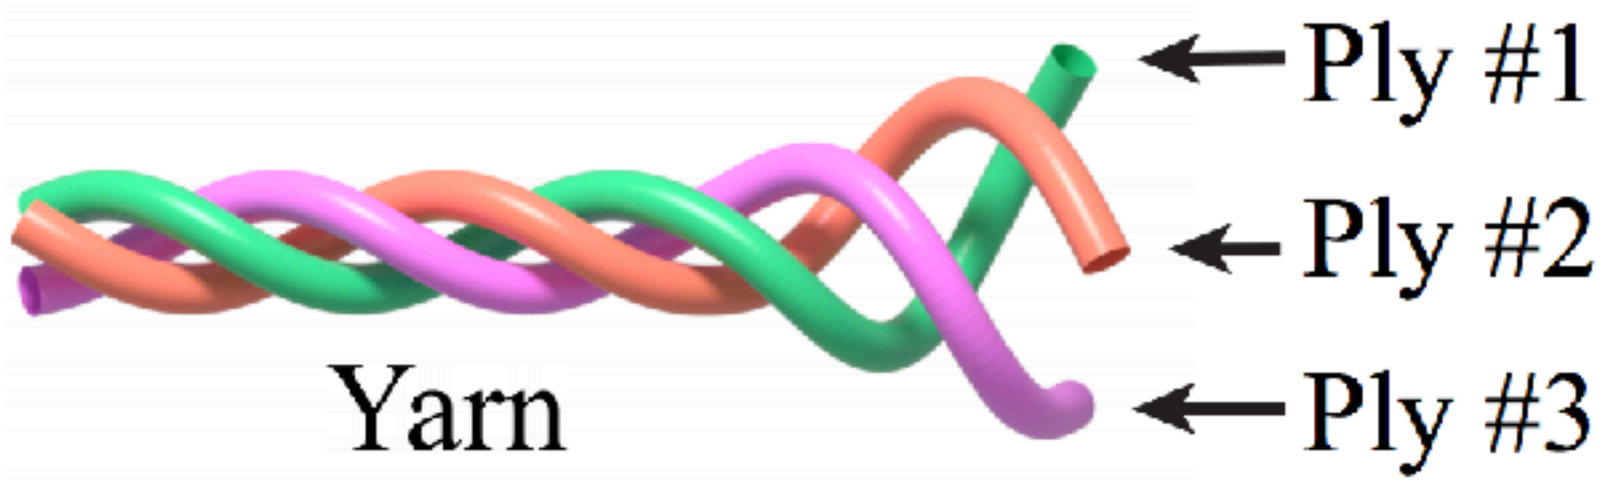

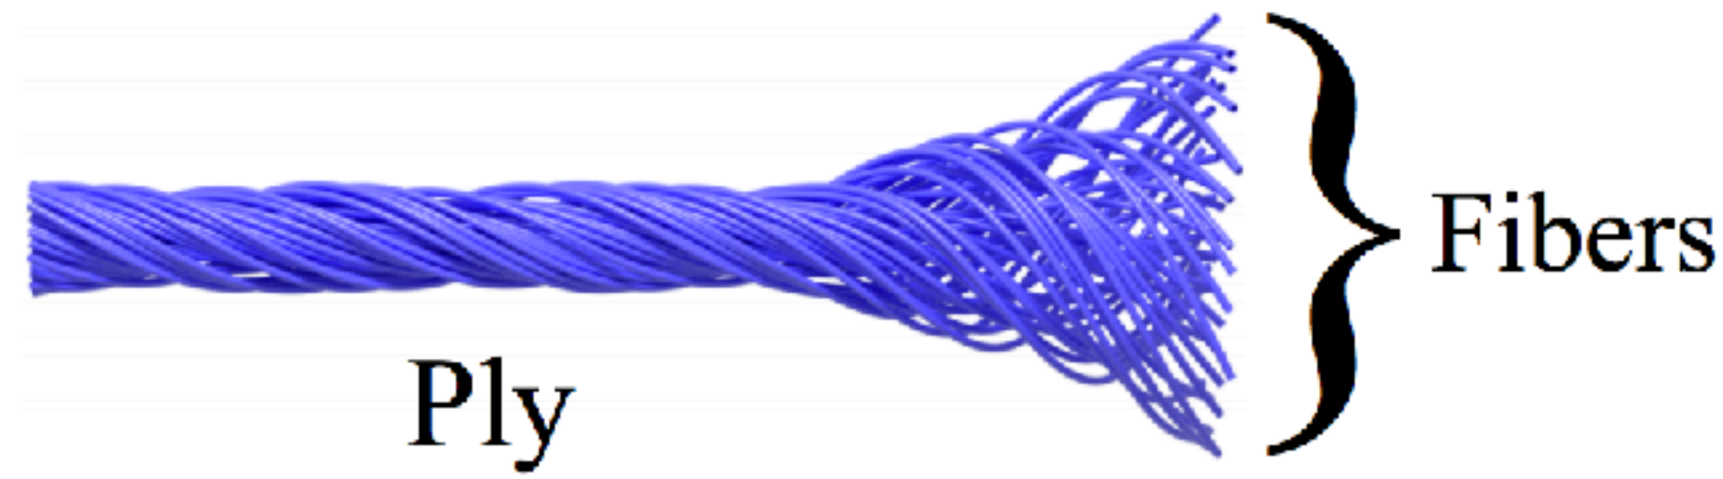

- Woven or knitted

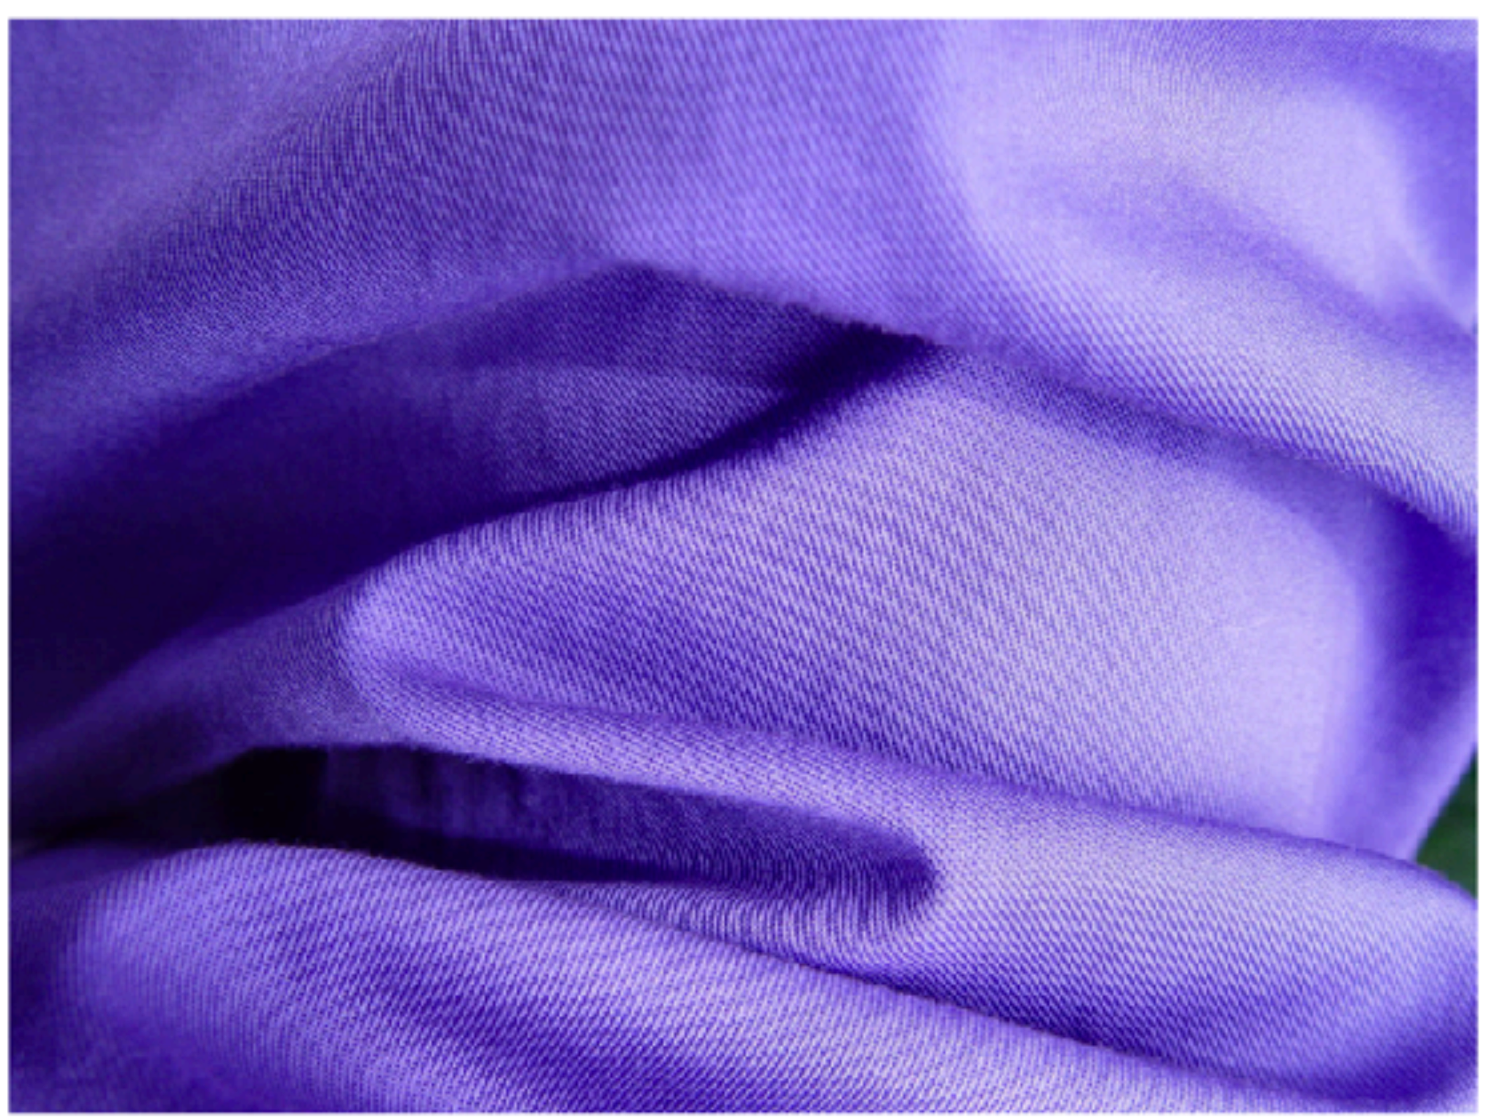

Cloth: Render as Surface

- Given the weaving pattern, calculate the overall behavior
- Render using a BRDF

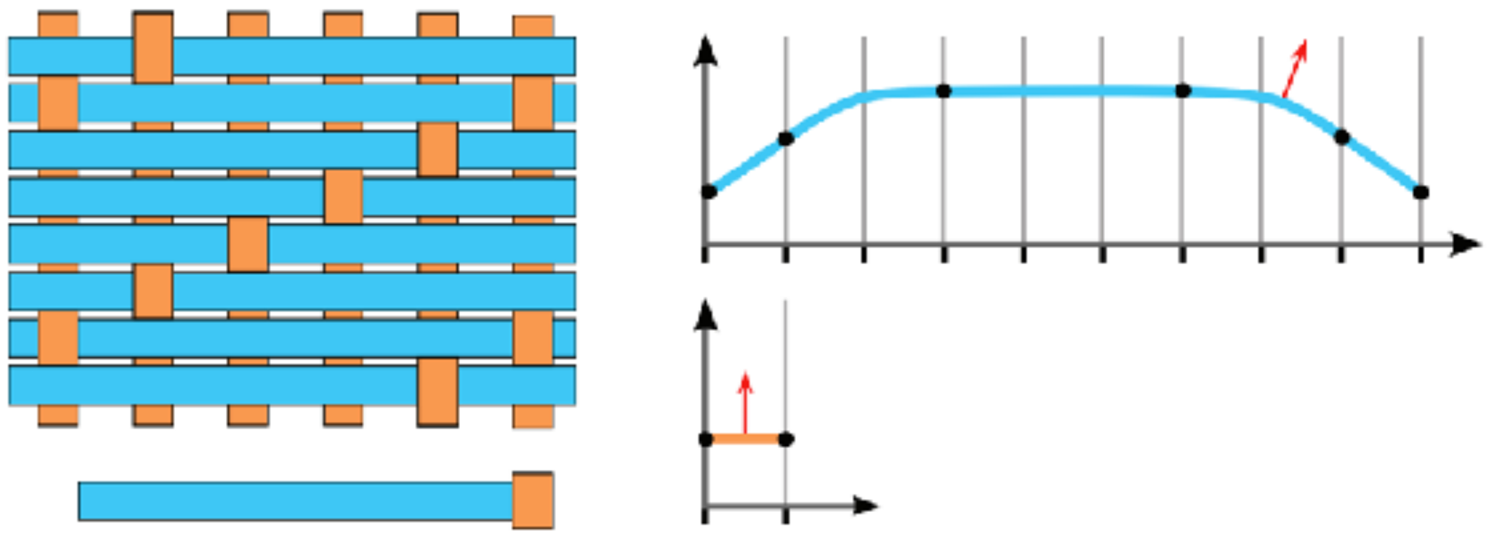

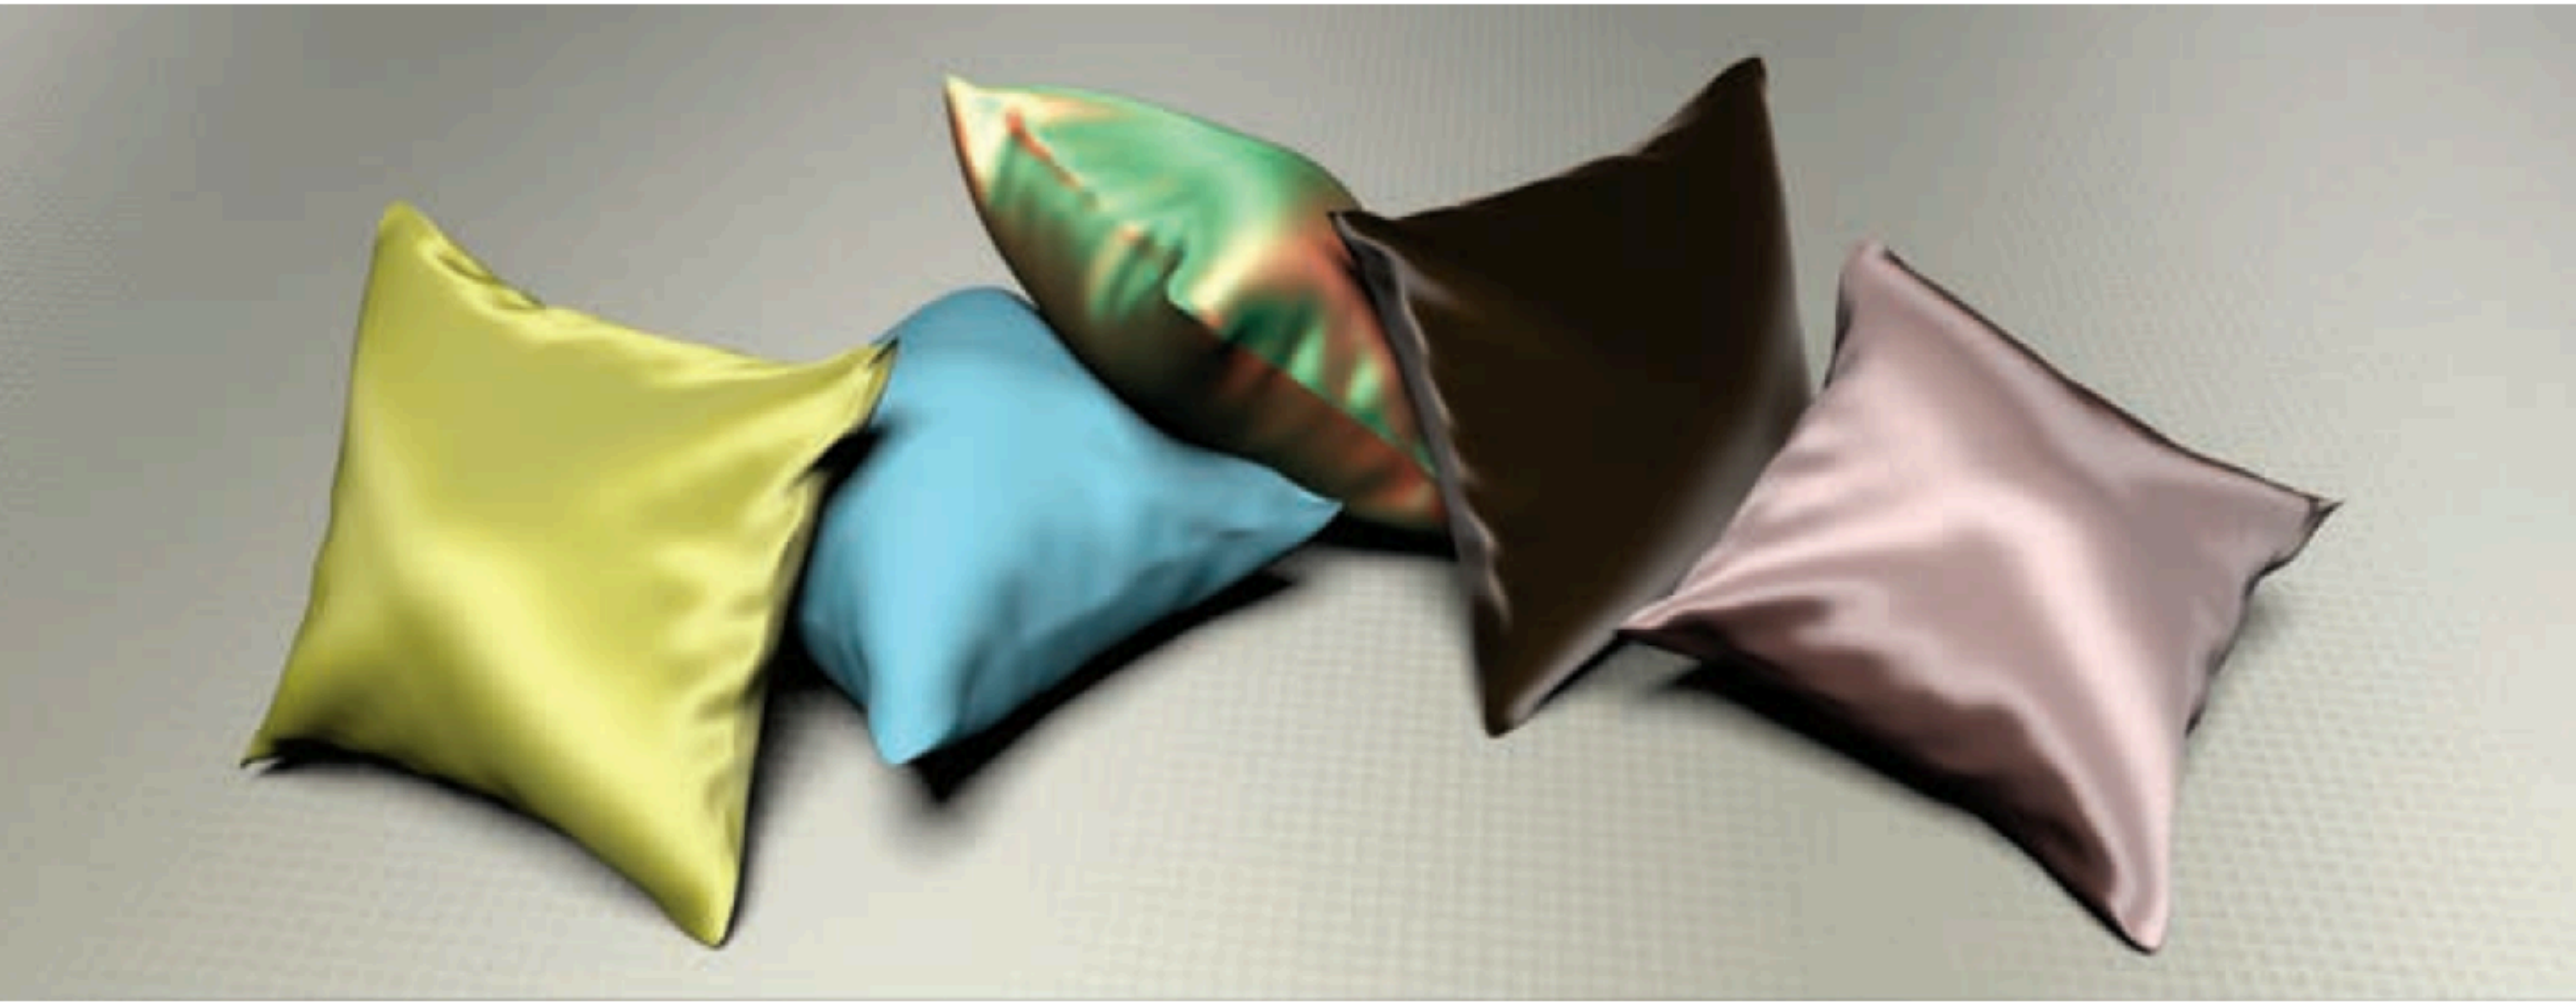

[Sadeghi et al. 2013]

Render as Surface — Limitation

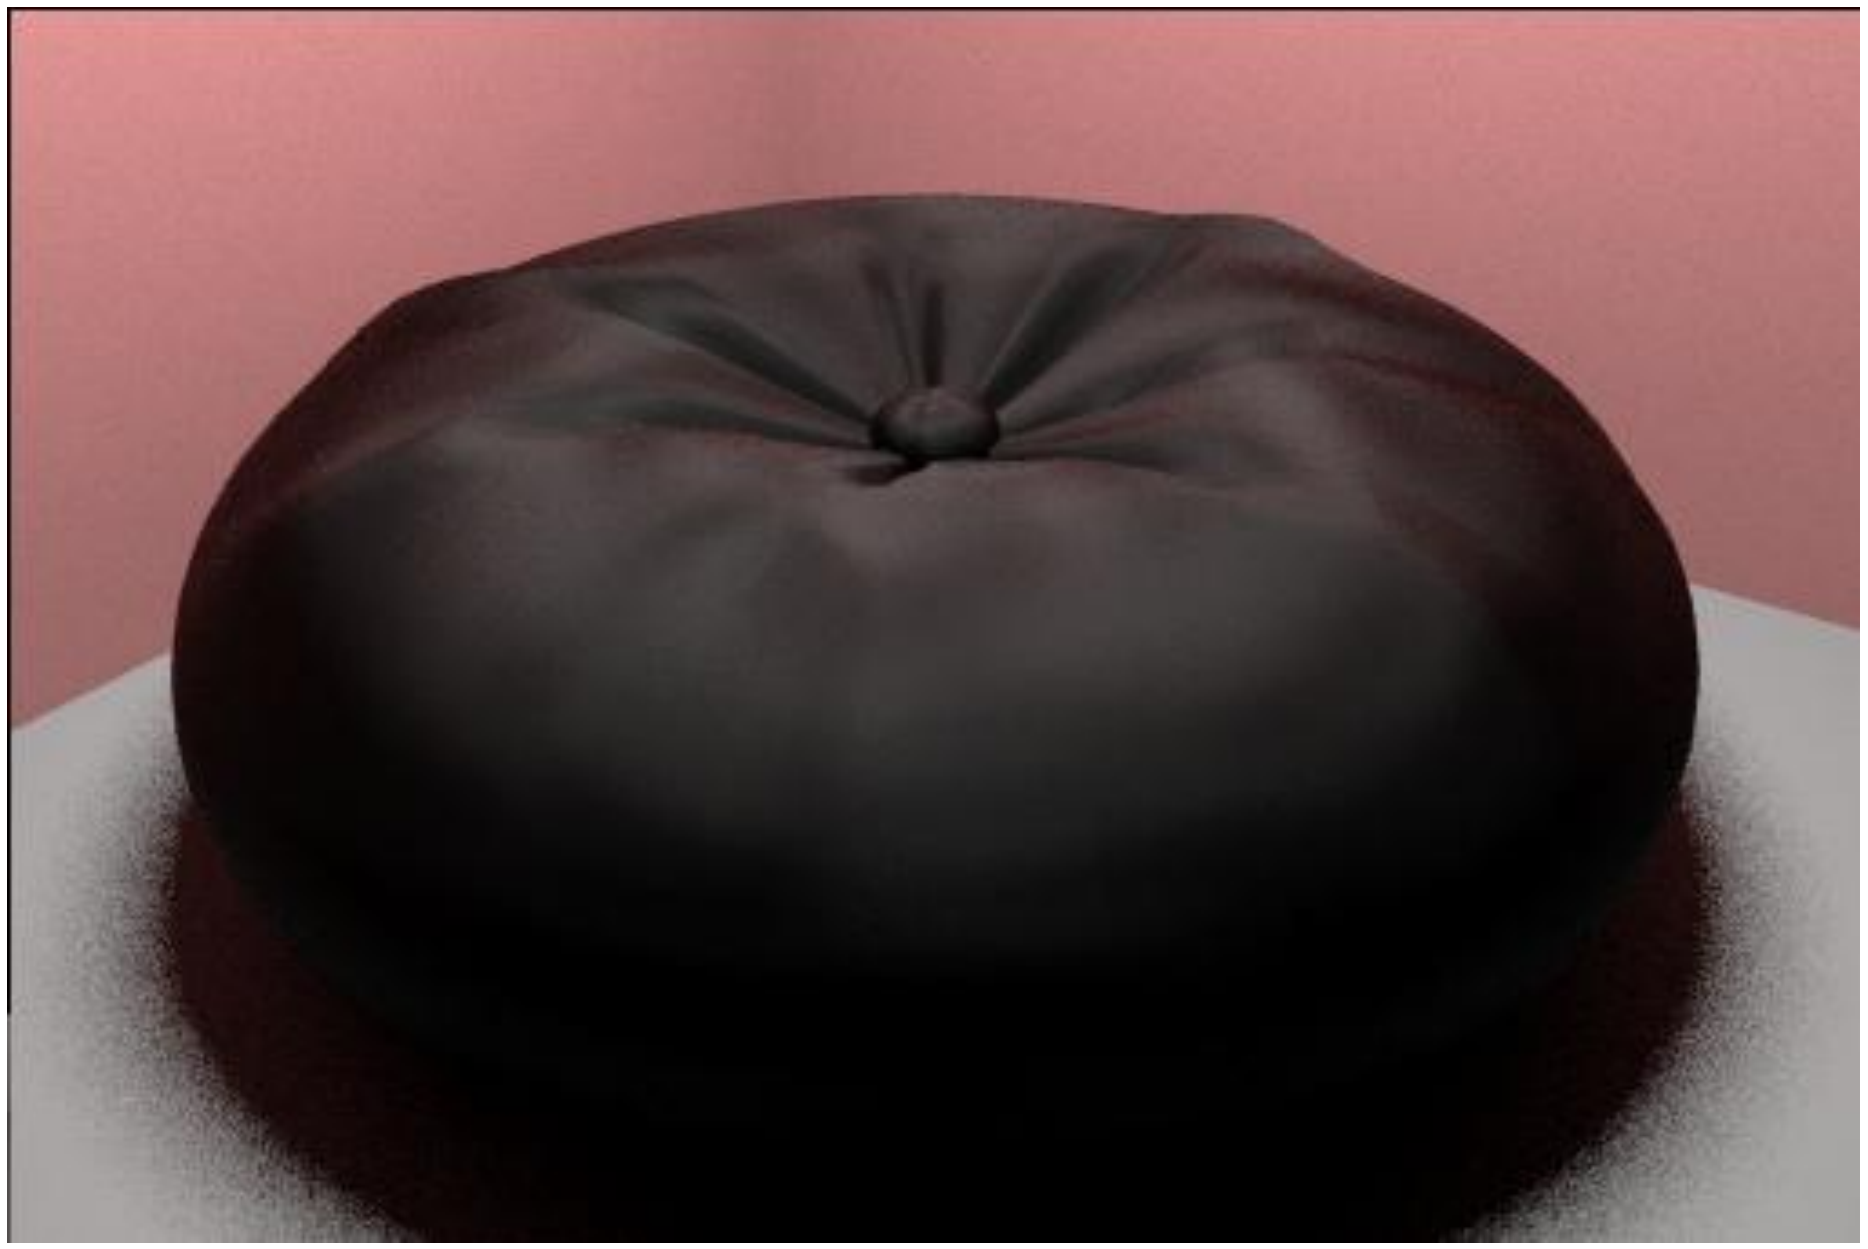

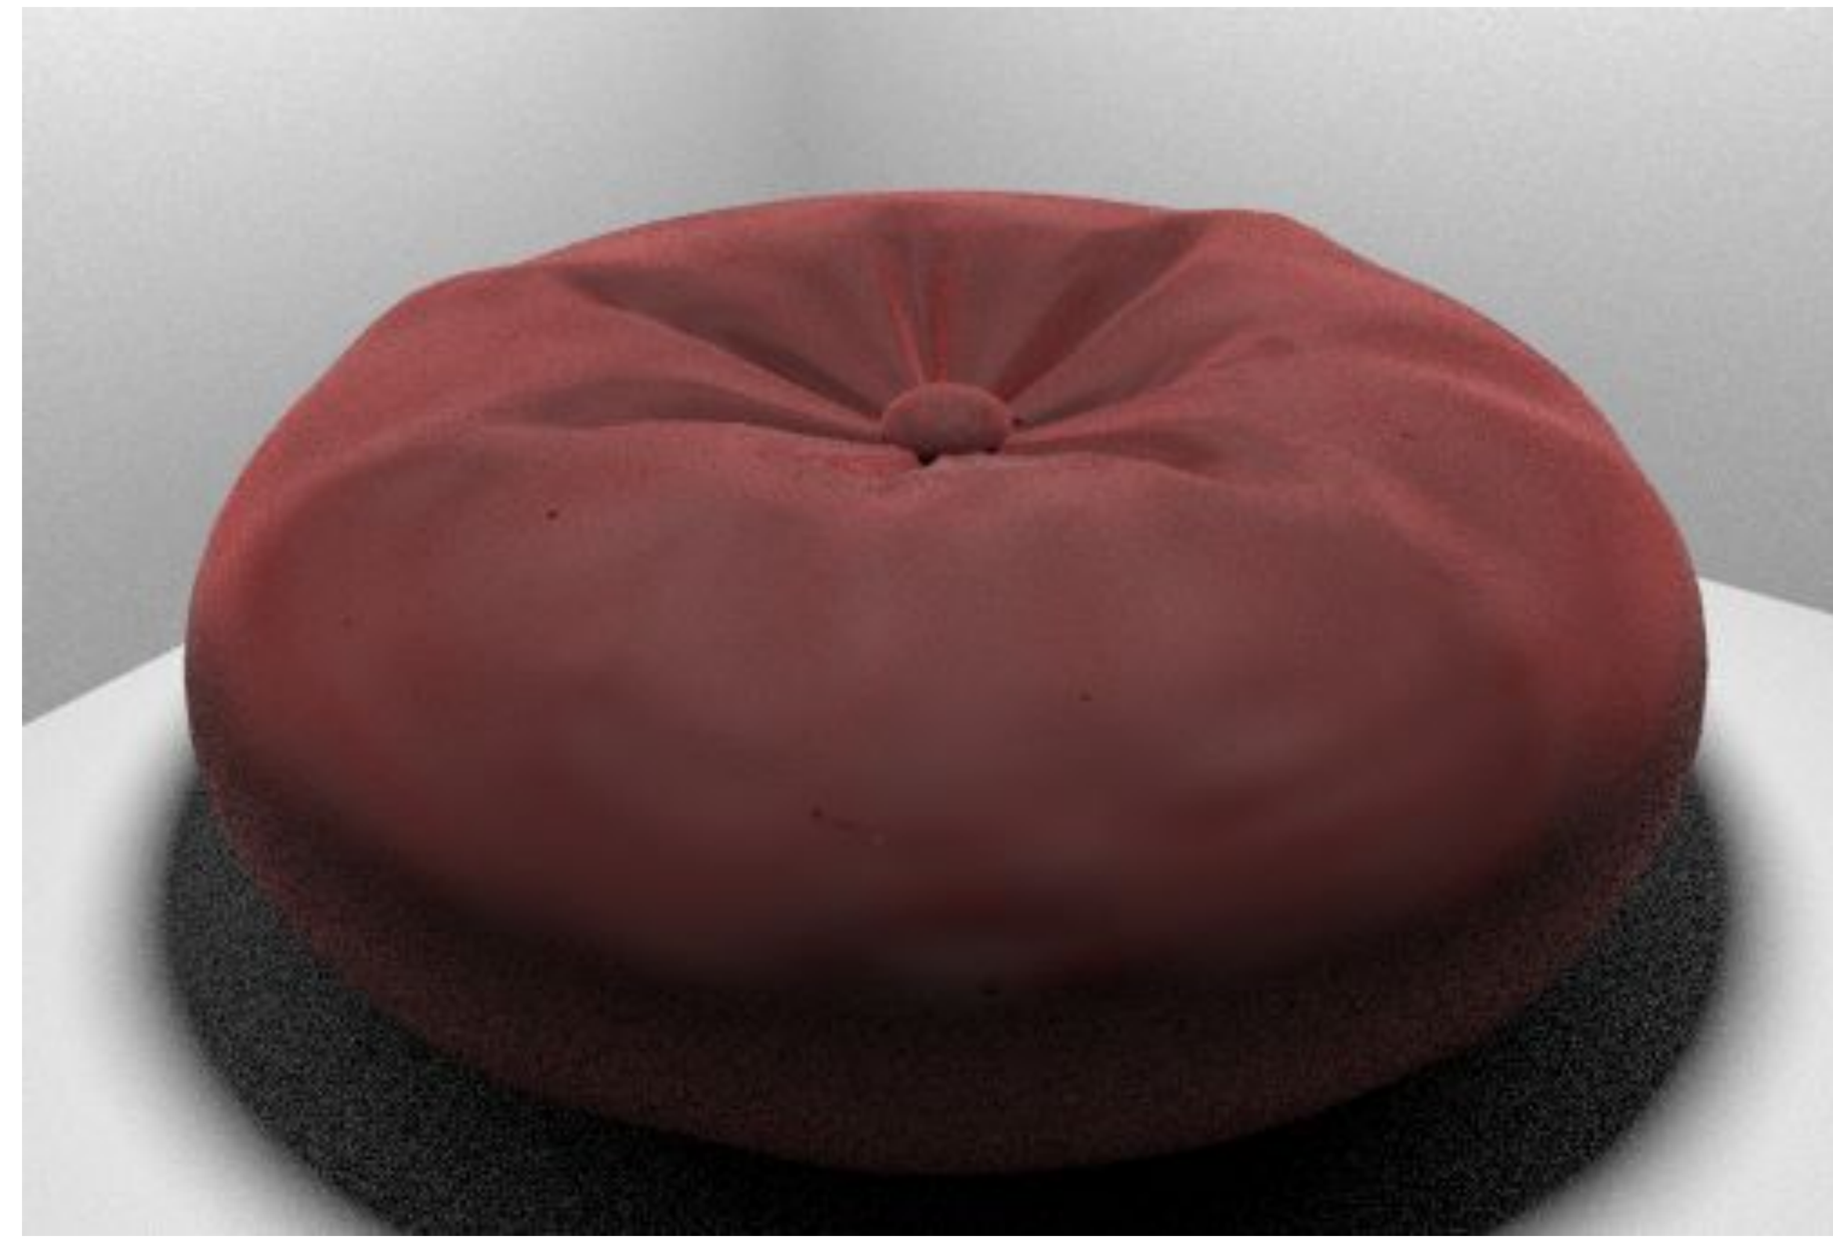

[Westin et al. 1992]

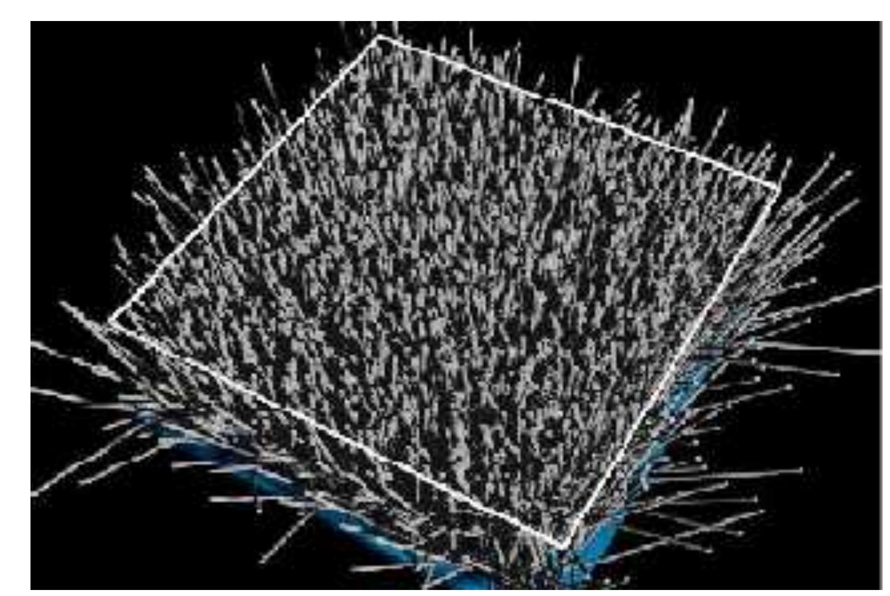

✓ ✗

Cloth: Render as Participating Media

- Properties of individual fibers & their distribution -> scattering parameters
- Render as a participating medium

[Jakob et al. 2010]

[Schroder et al. 2011]

Cloth: Render as Actual Fibers

- Render every fiber explicitly!

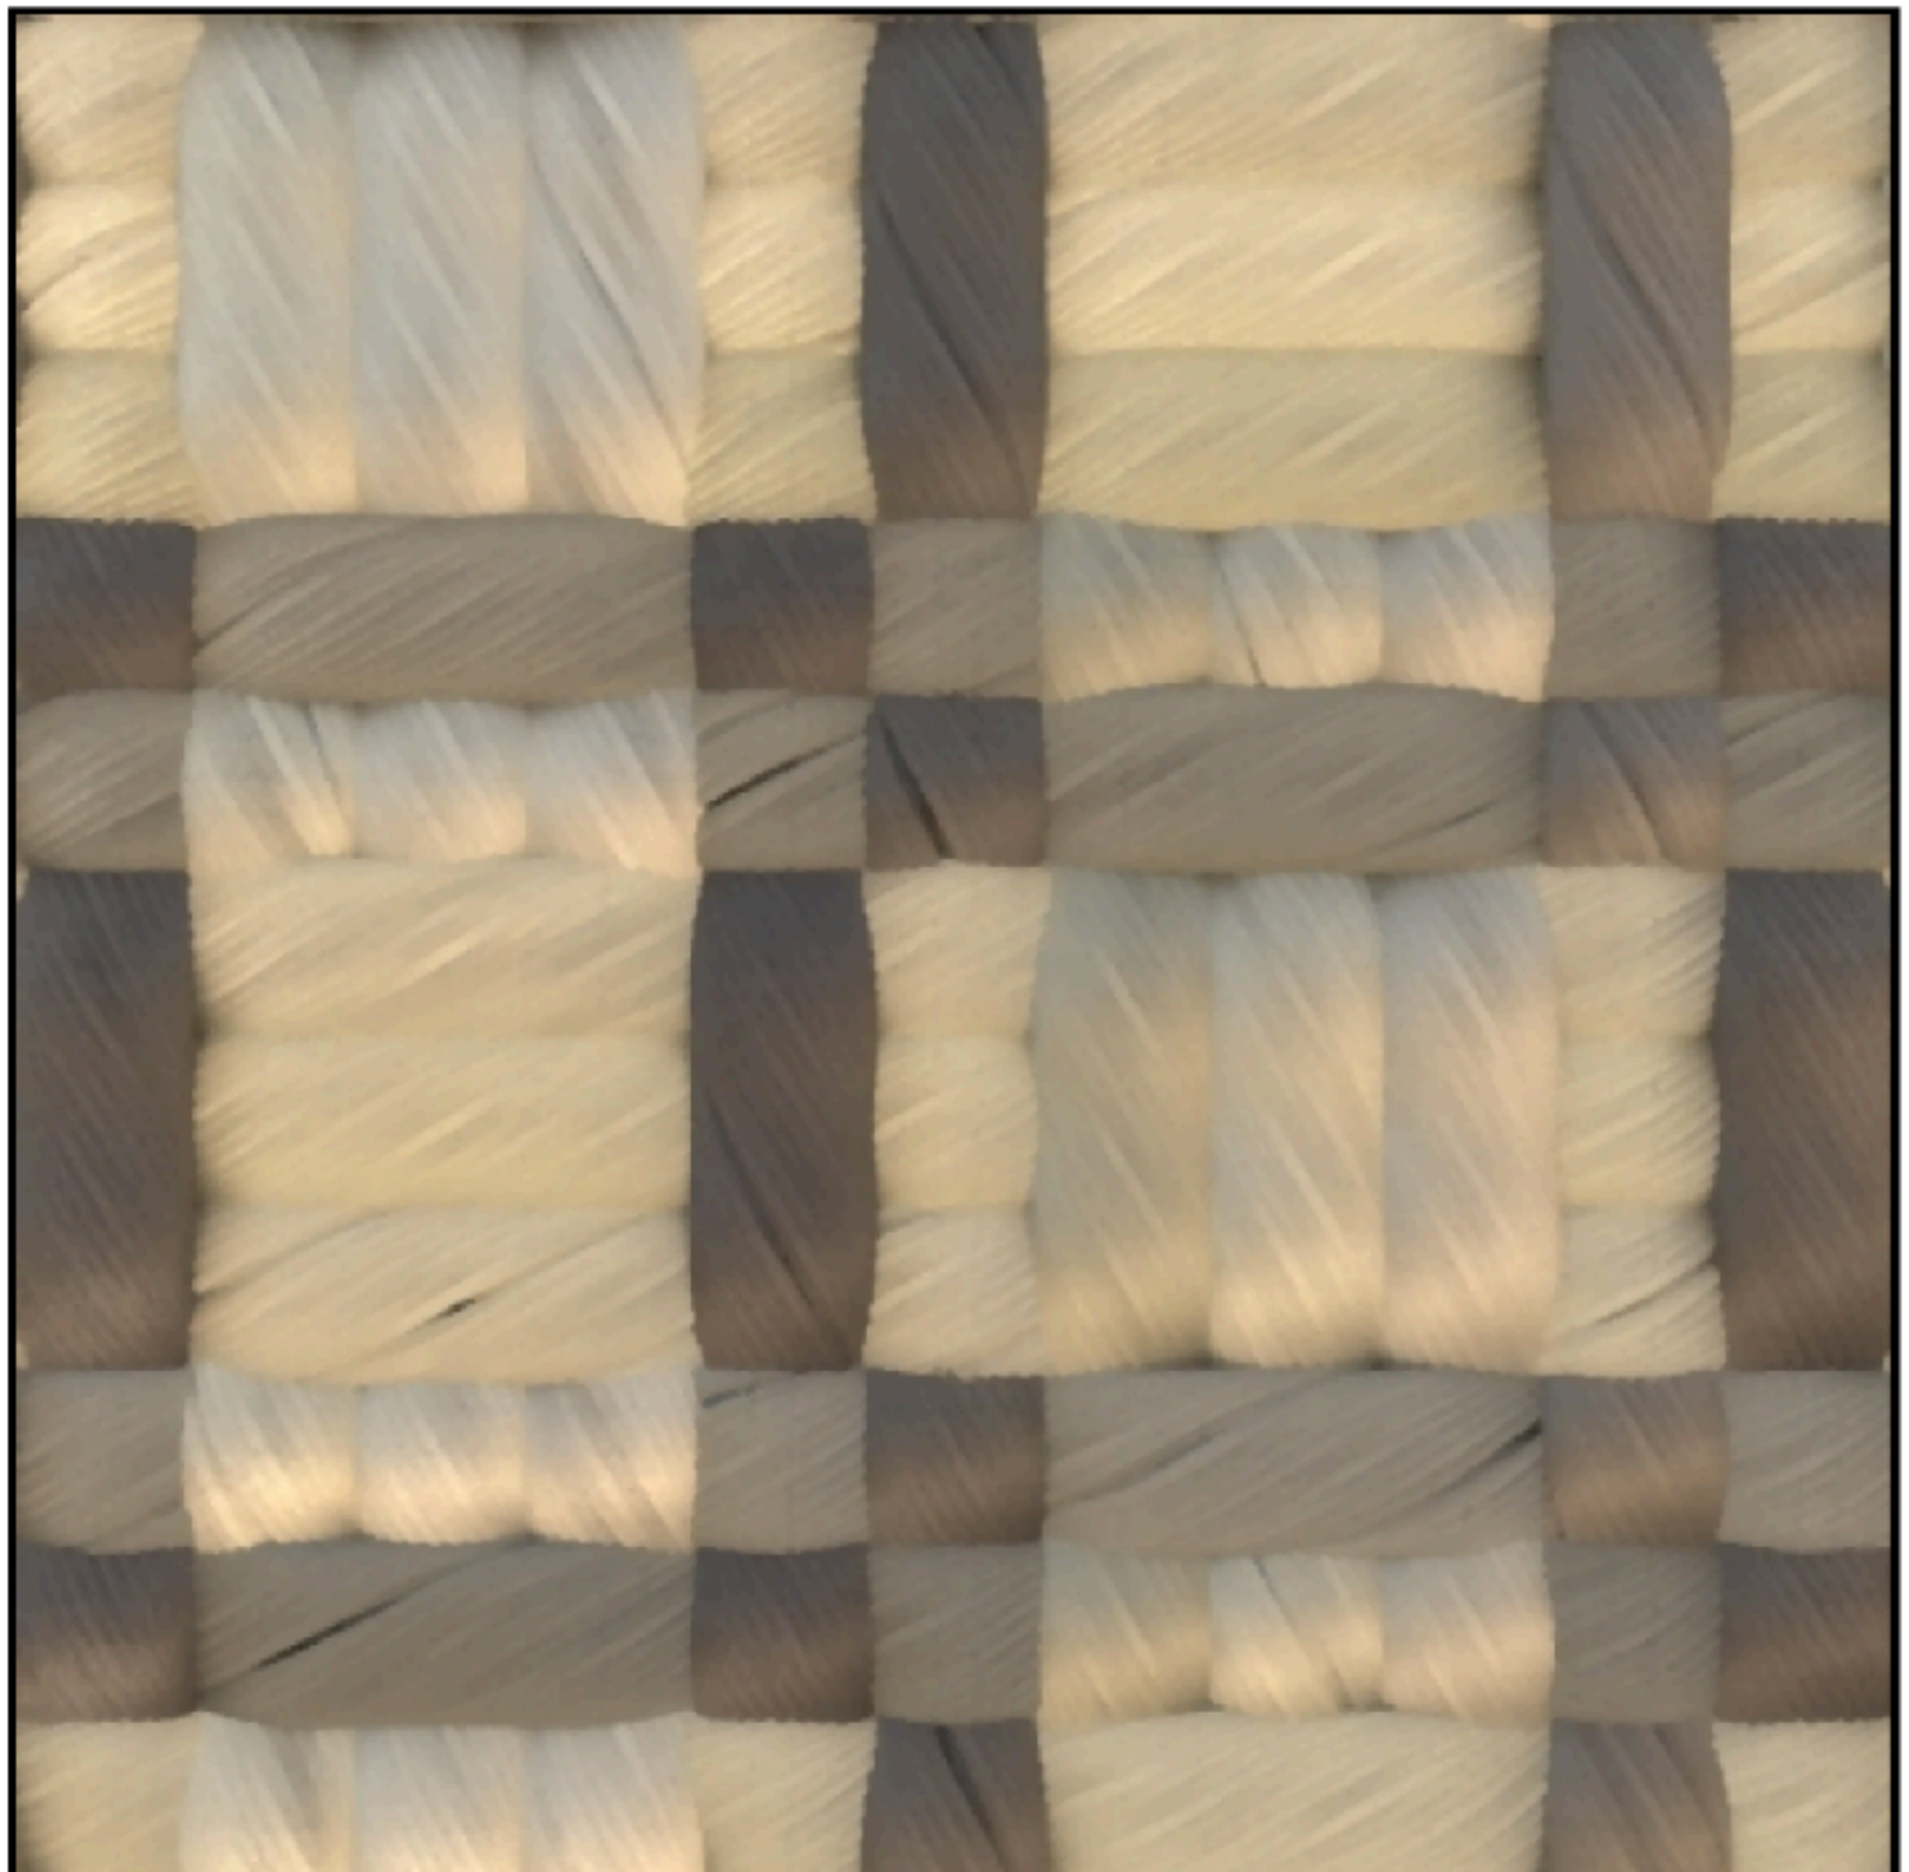

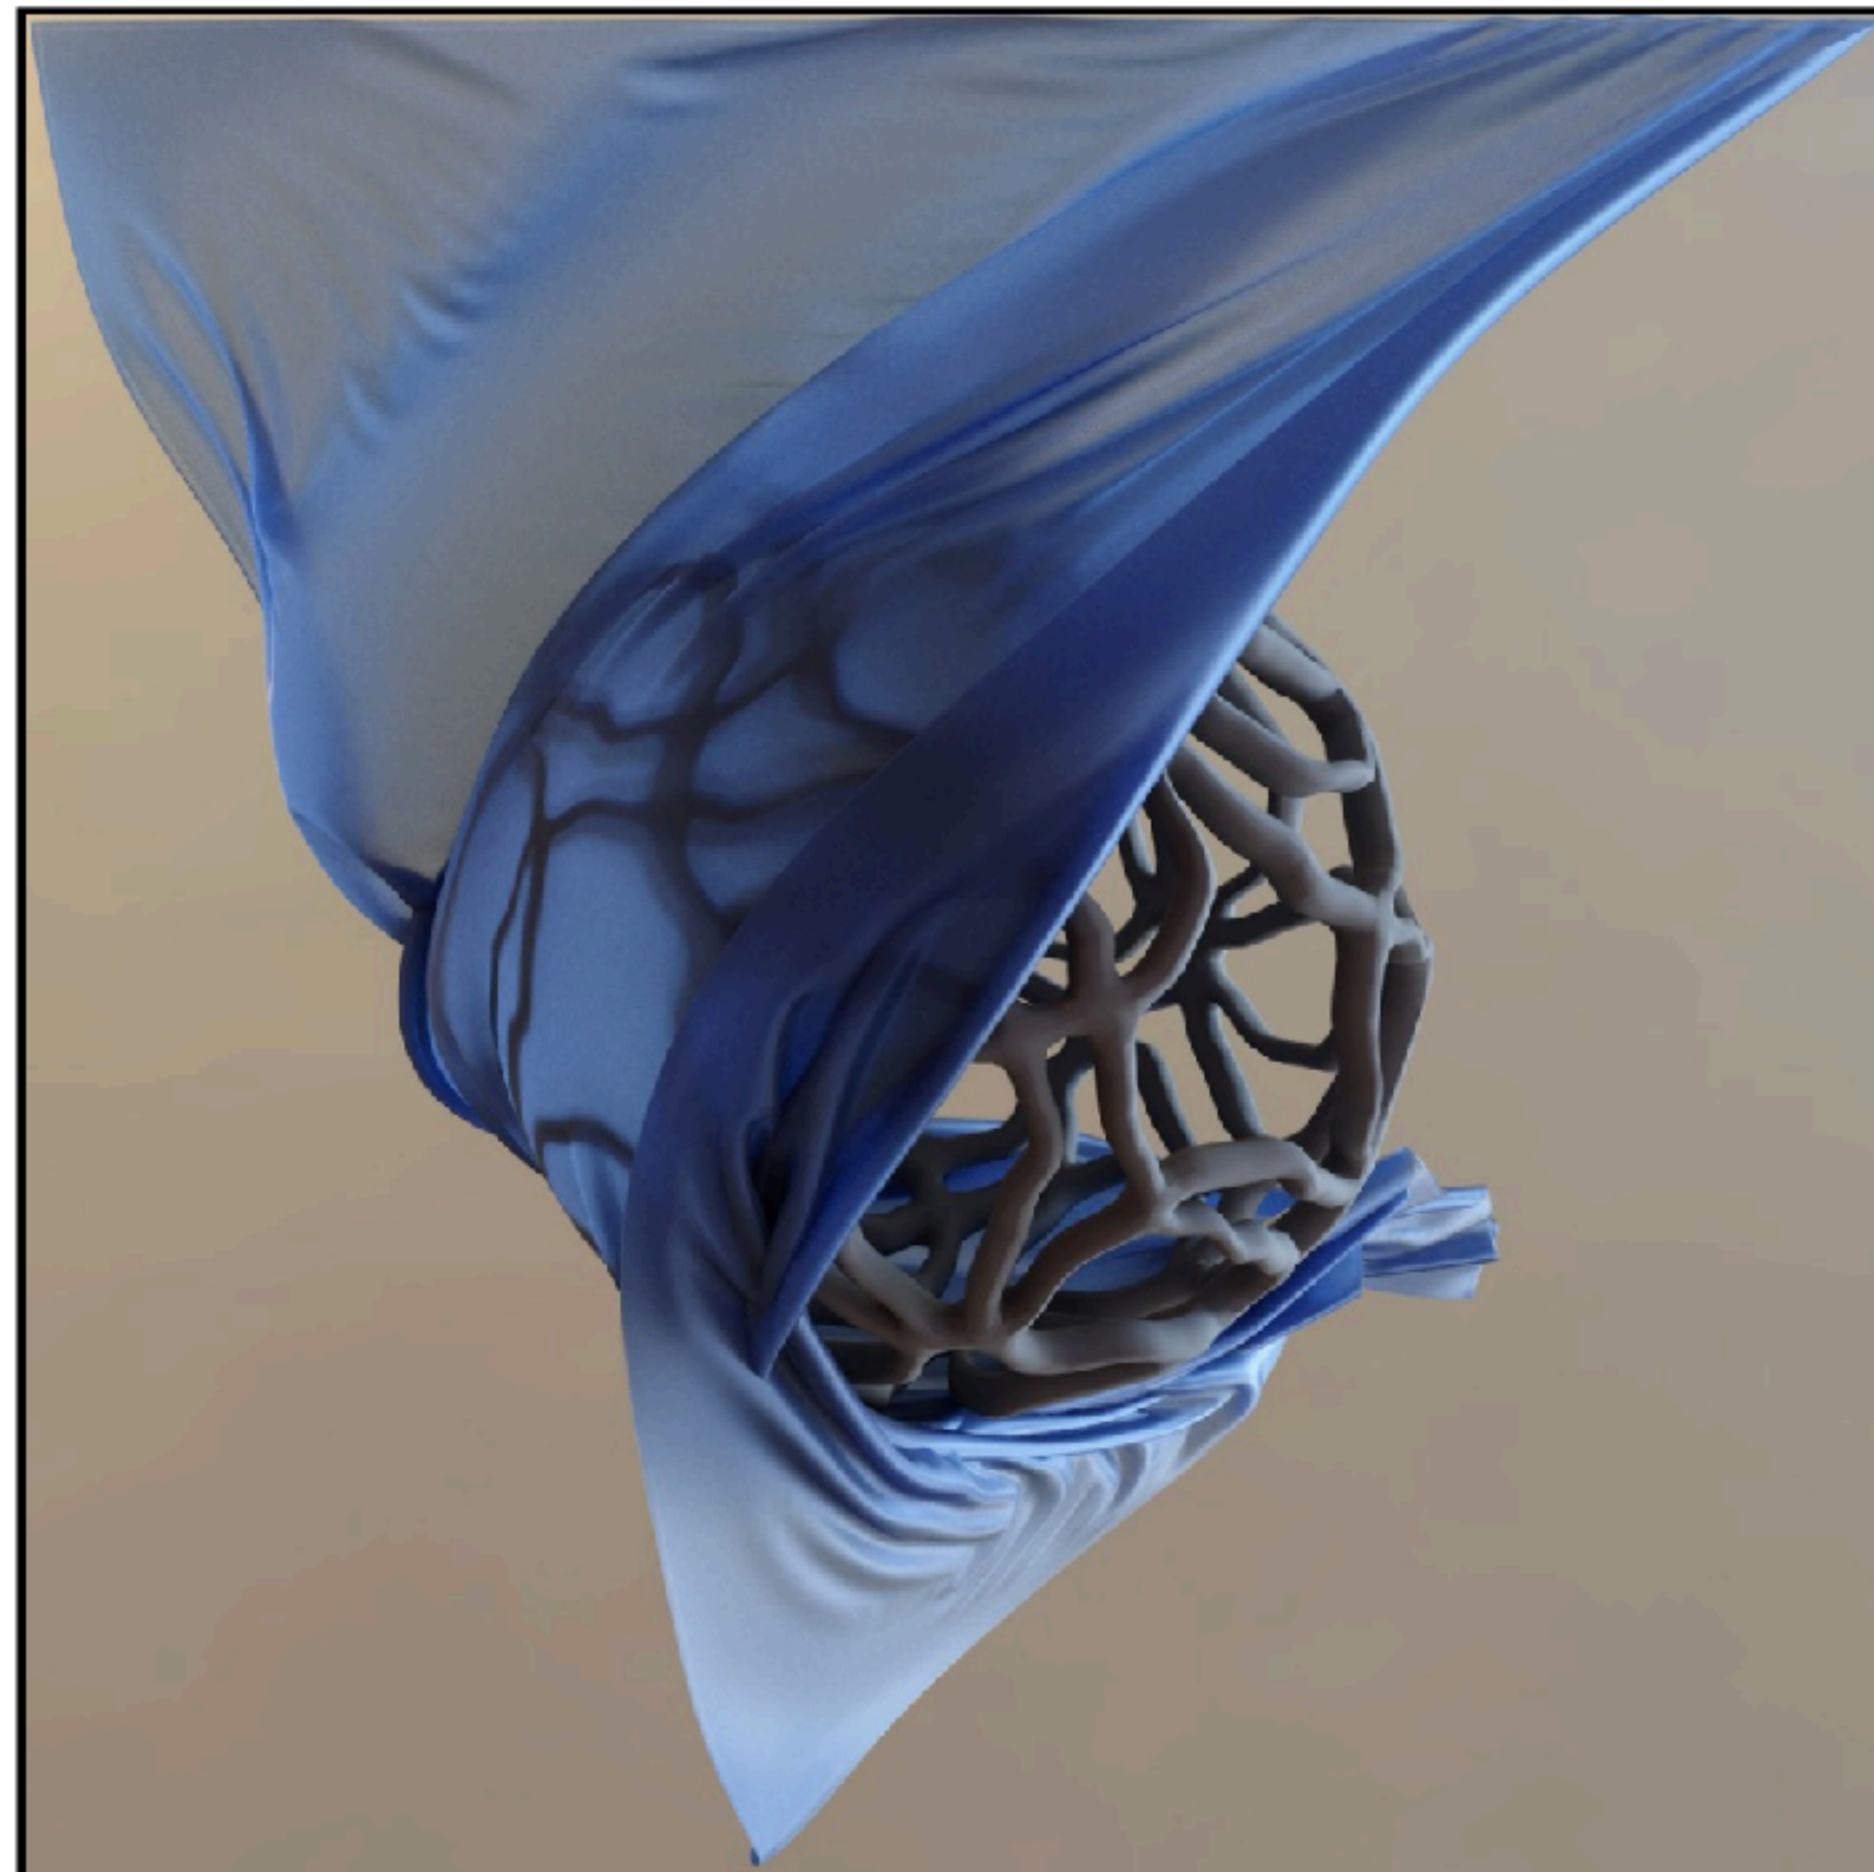

[Kai Schroder]

Cloth: Demo

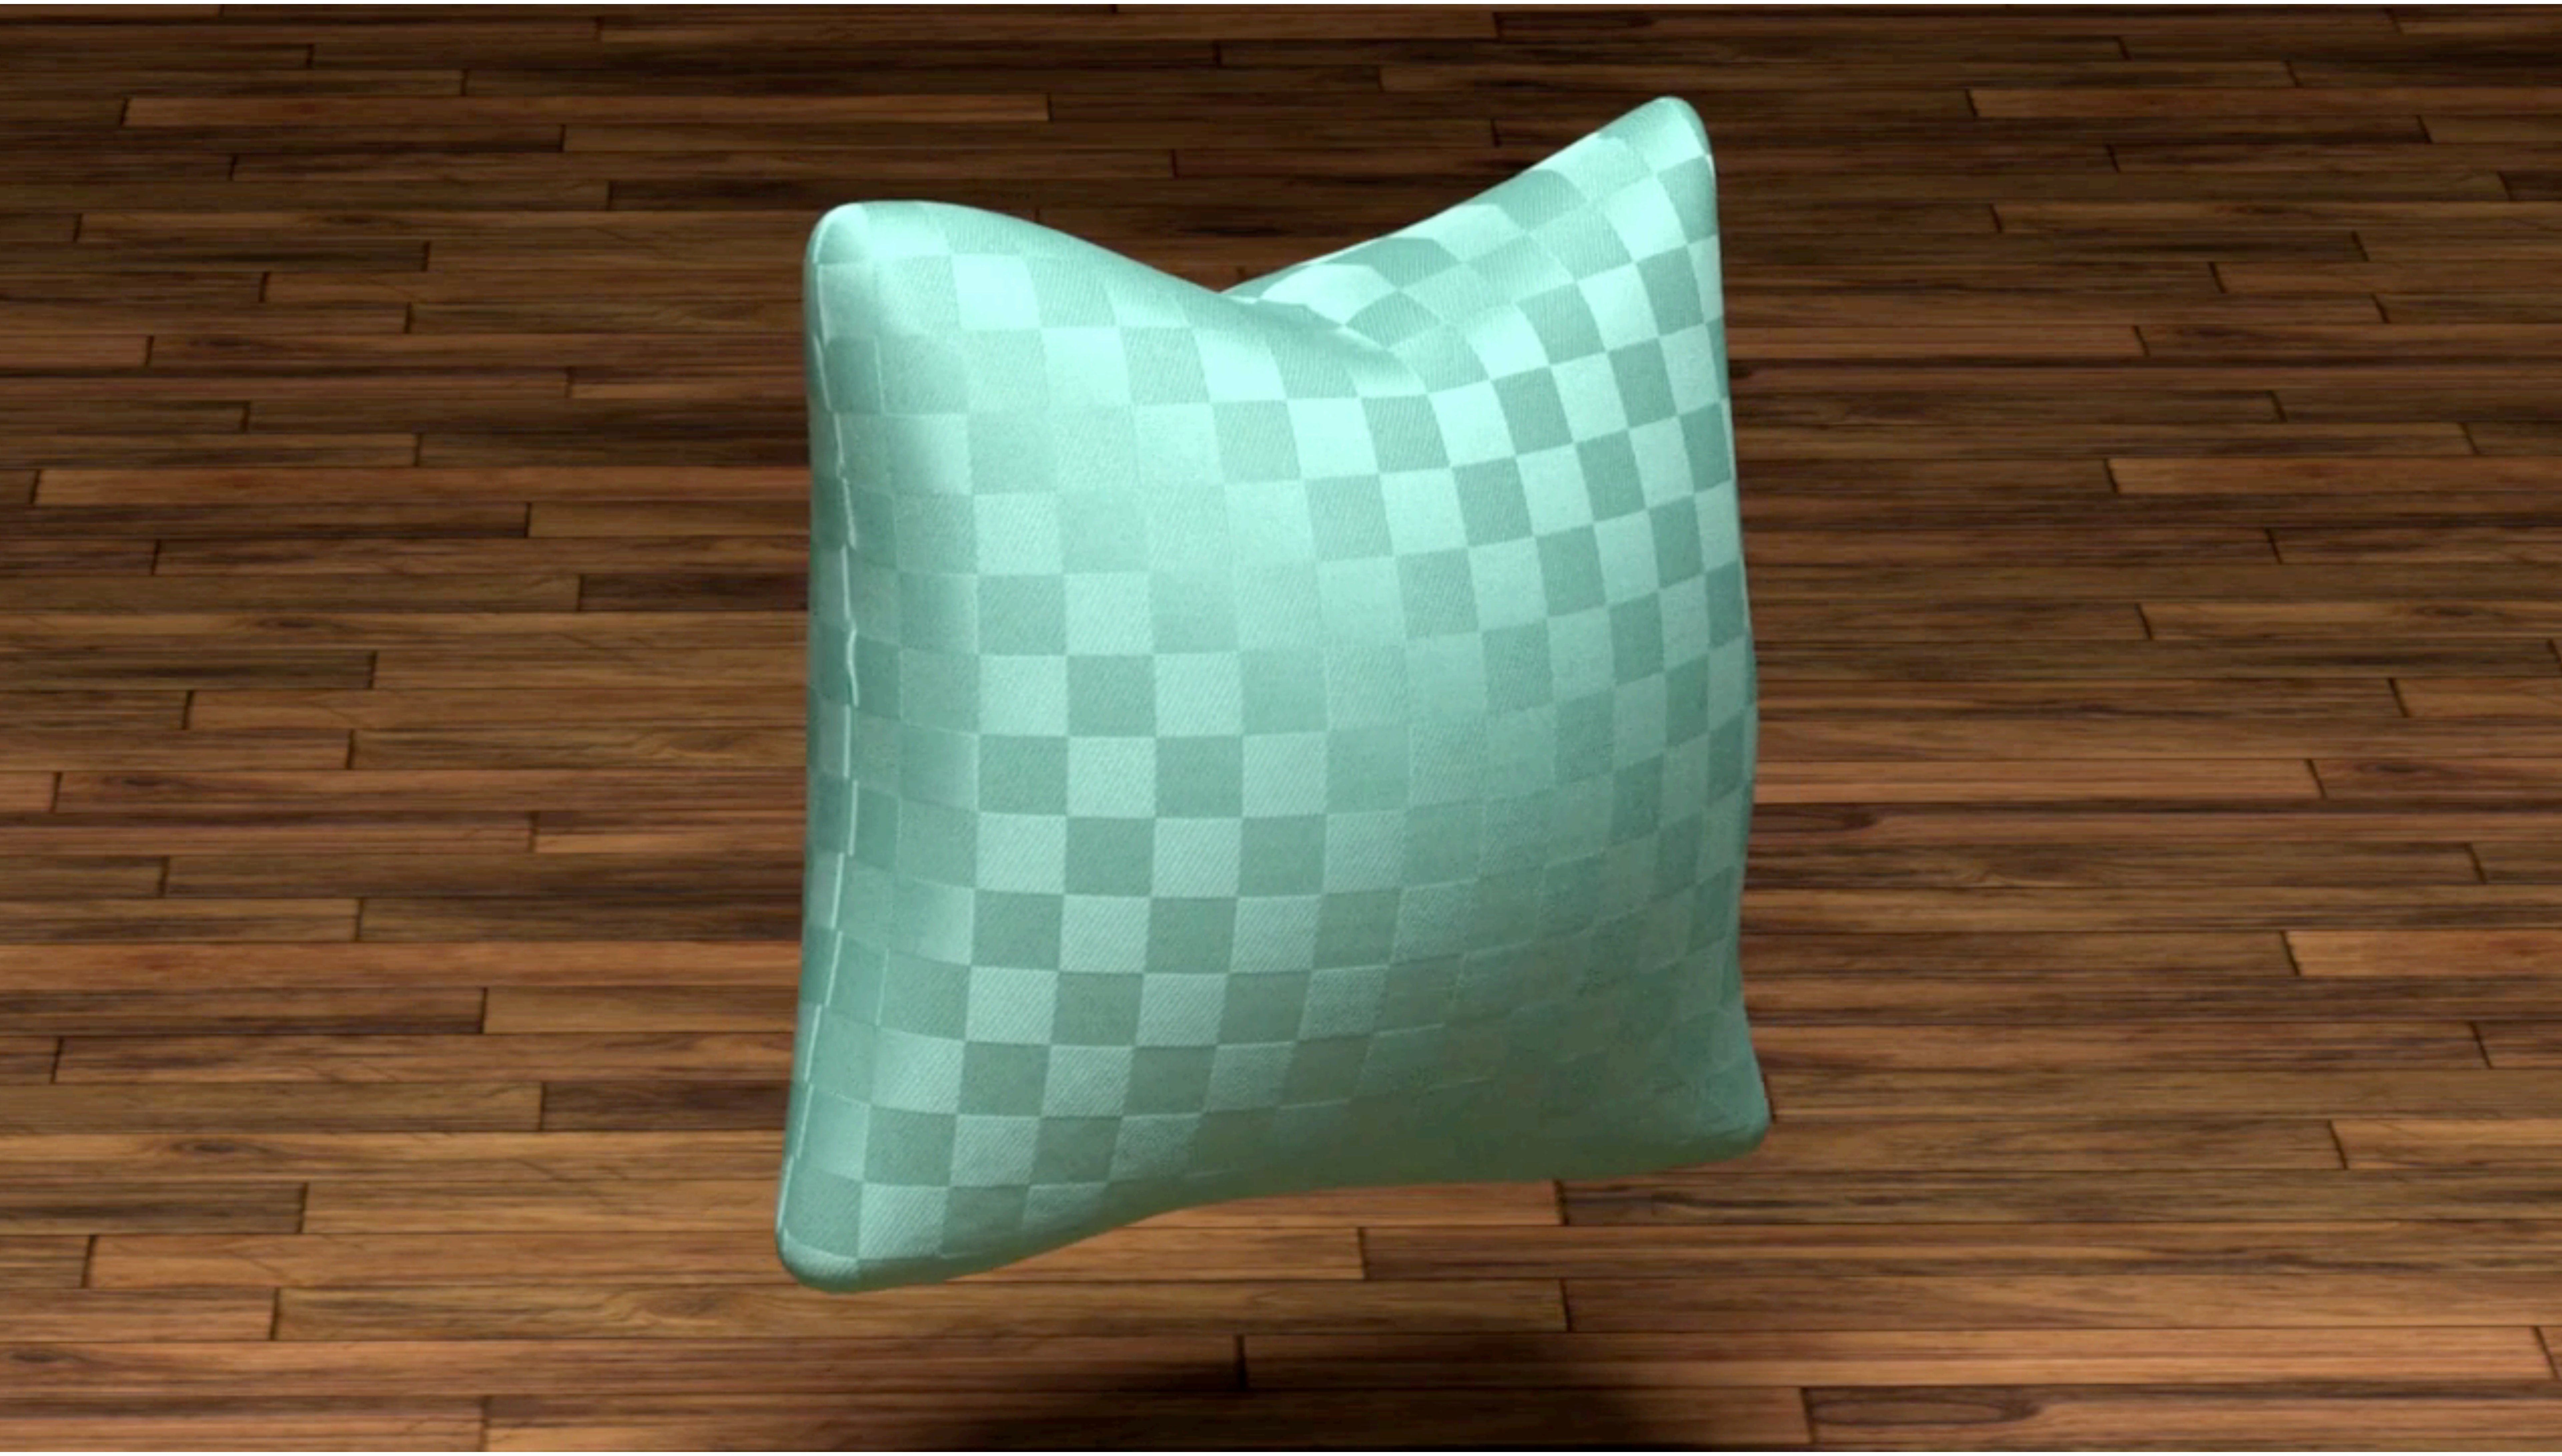

[Shuang et al. 2012]

Cloth: Application

[The BFG. 2016 Disney]

Detailed Appearance: Motivation

- Not looking realistic, why?

[Car rendered in NVIDIA Iray]

[Mouse rendered in Autodesk 3DS Max]

Real world is more complicated

[Real photograph of a car]

[Real video of a mouse]

Why details?

Microfacet
model

Why details?

[Yan et al.
2014, 2016]

Why details?

[Yan et al. 2014, 2016]

Recap: Microfacet BRDF

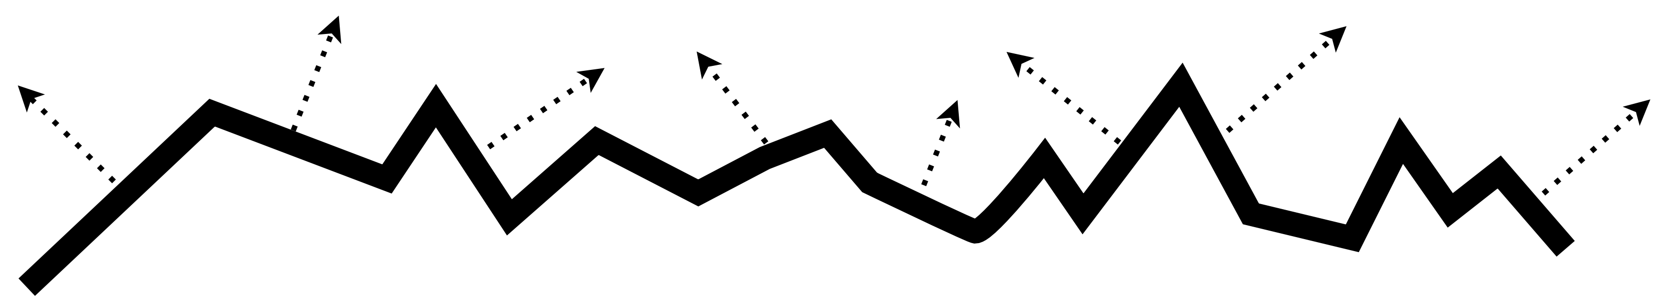

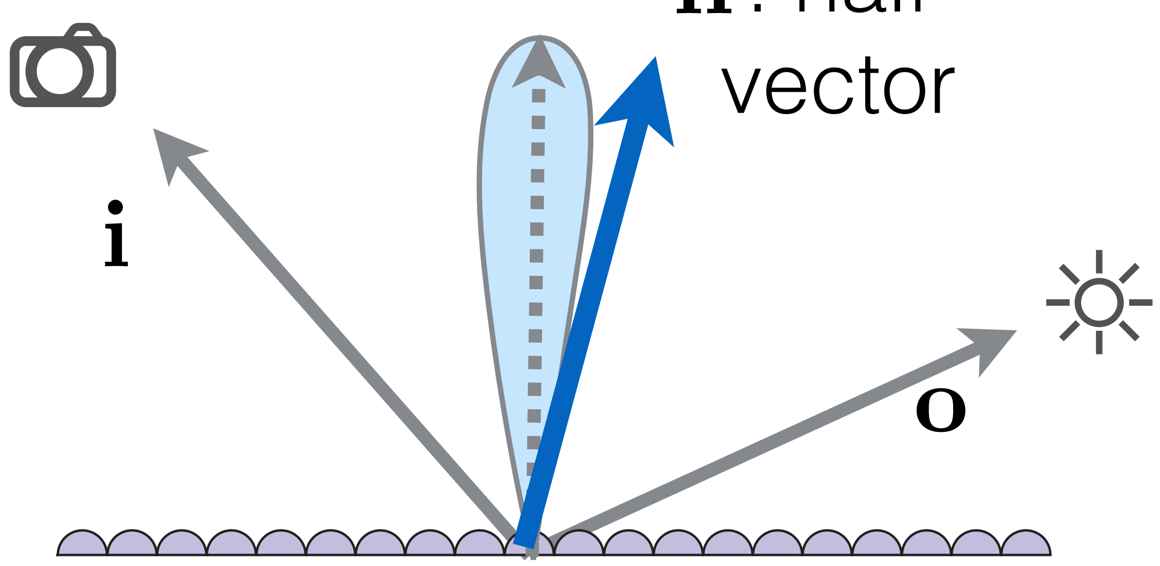

Surface = Specular microfacets + statistical normals

$$f(\mathbf{i}, \mathbf{o}) = \frac{\mathbf{F}(\mathbf{i}, \mathbf{h})\mathbf{G}(\mathbf{i}, \mathbf{o}, \mathbf{h})\boxed{\mathbf{D}(\mathbf{h})}}{4(\mathbf{n}, \mathbf{i})(\mathbf{n}, \mathbf{o})}$$

NDF: Normal Distribution Function

Statistical NDF vs. Actual NDF

Distribution of Normals (NDF)

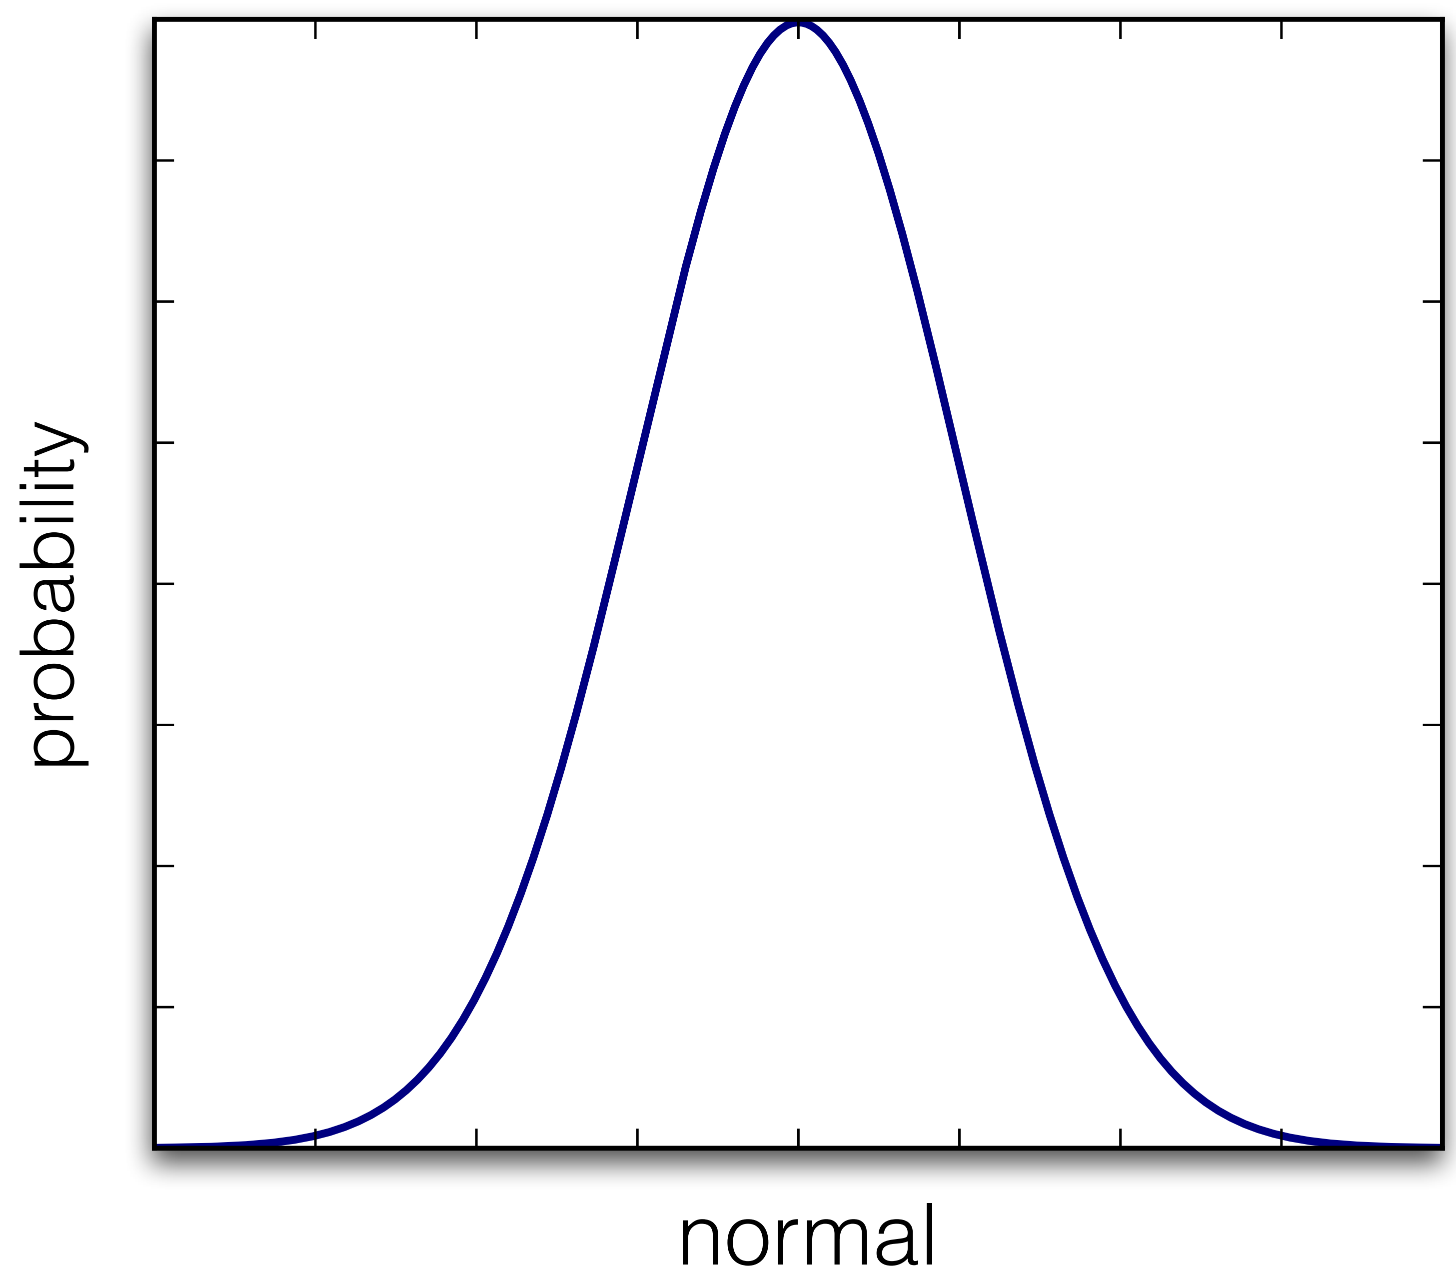

What we have
(microfacet — statistical)

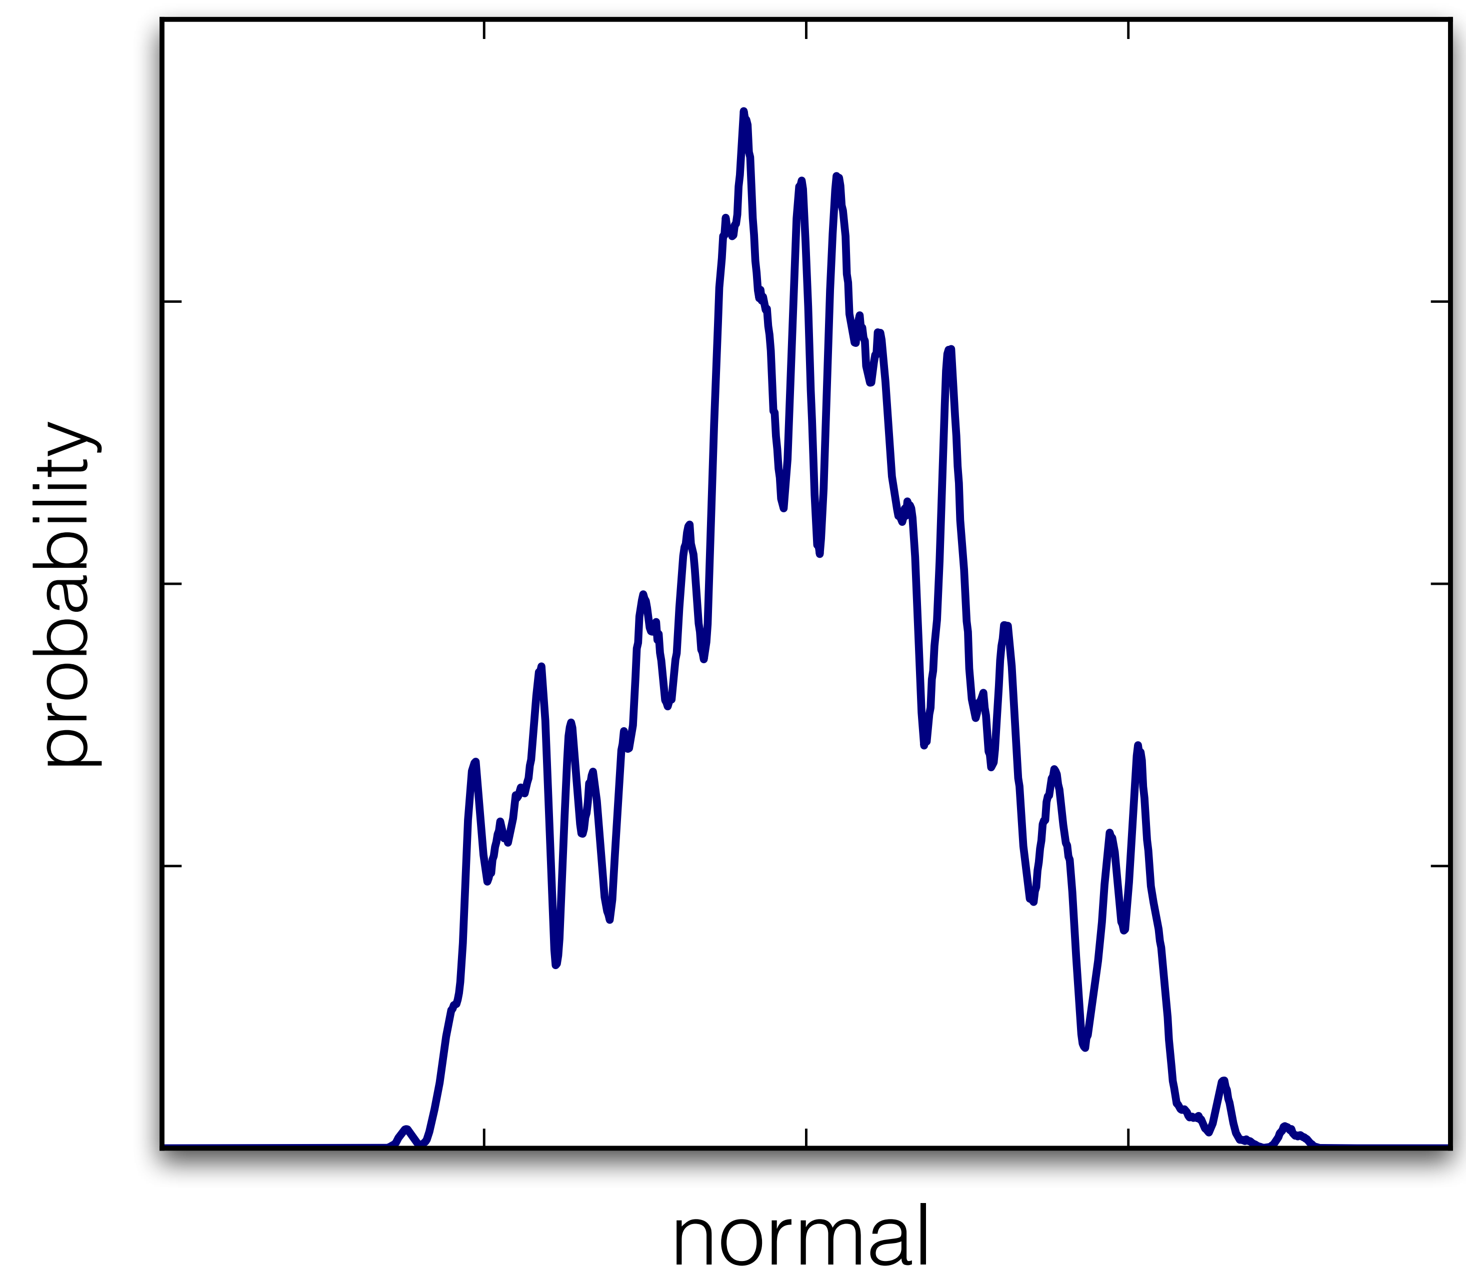

What we want

Define details

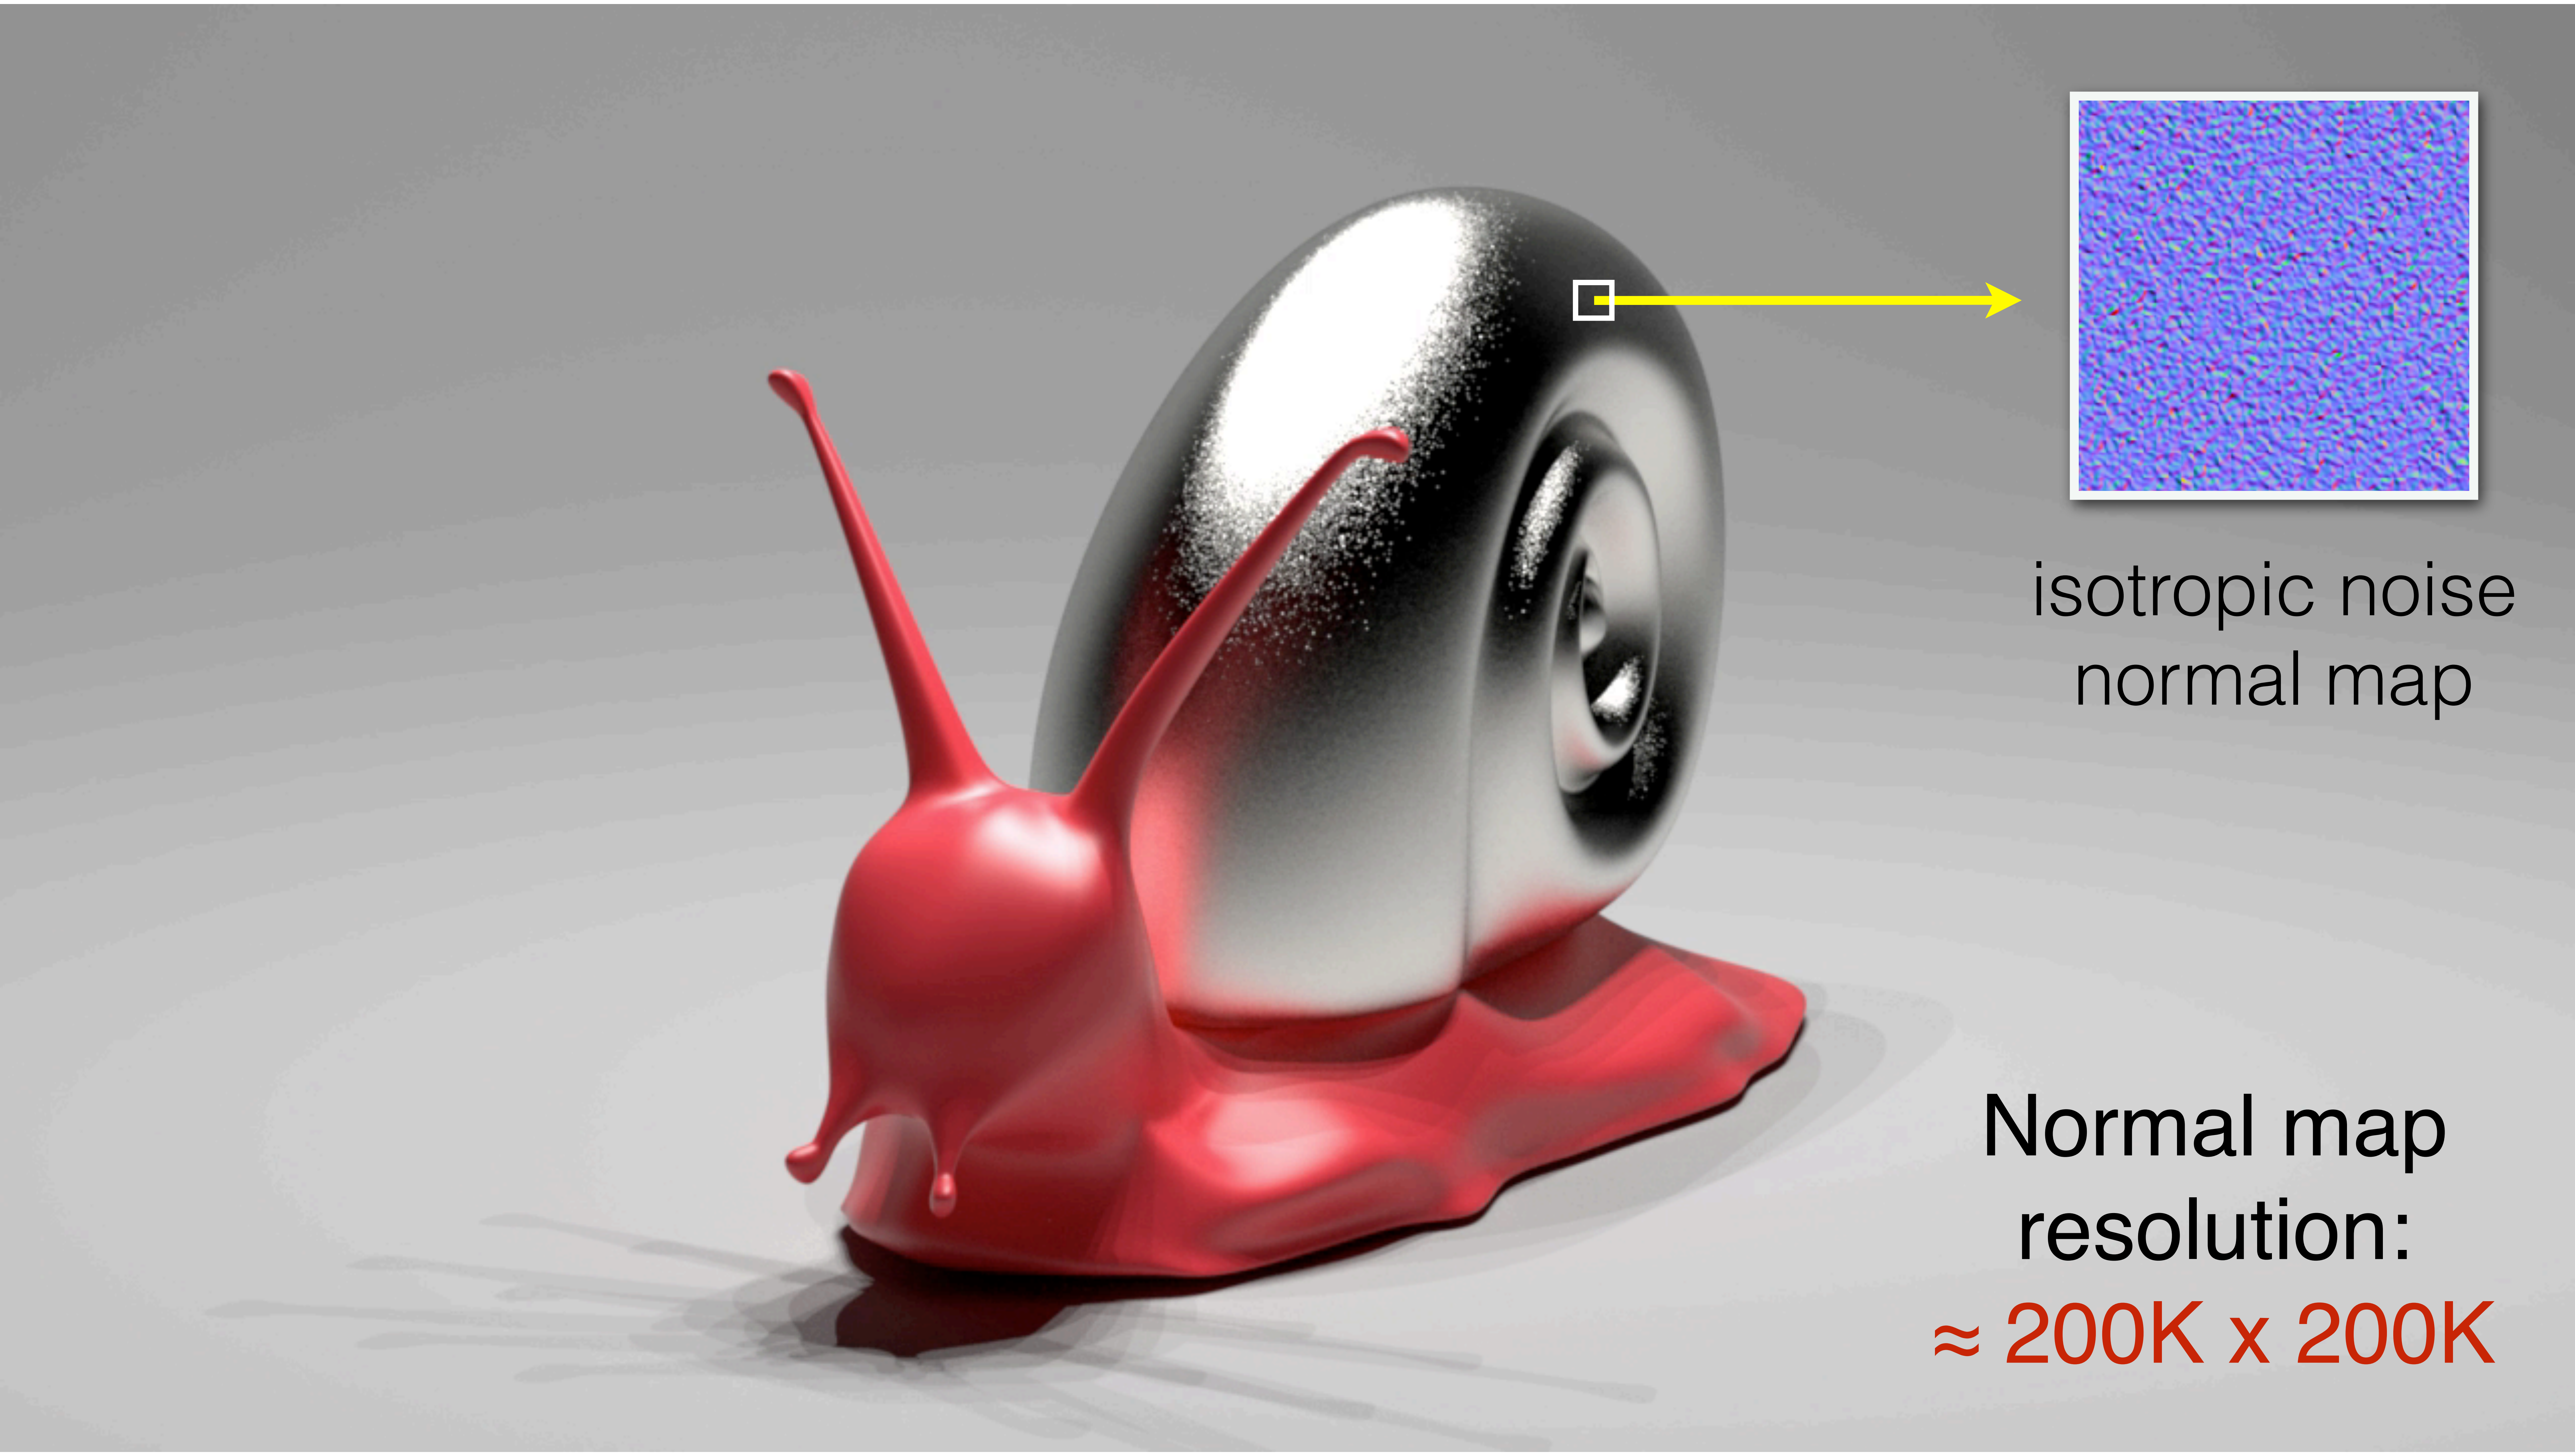

Define details

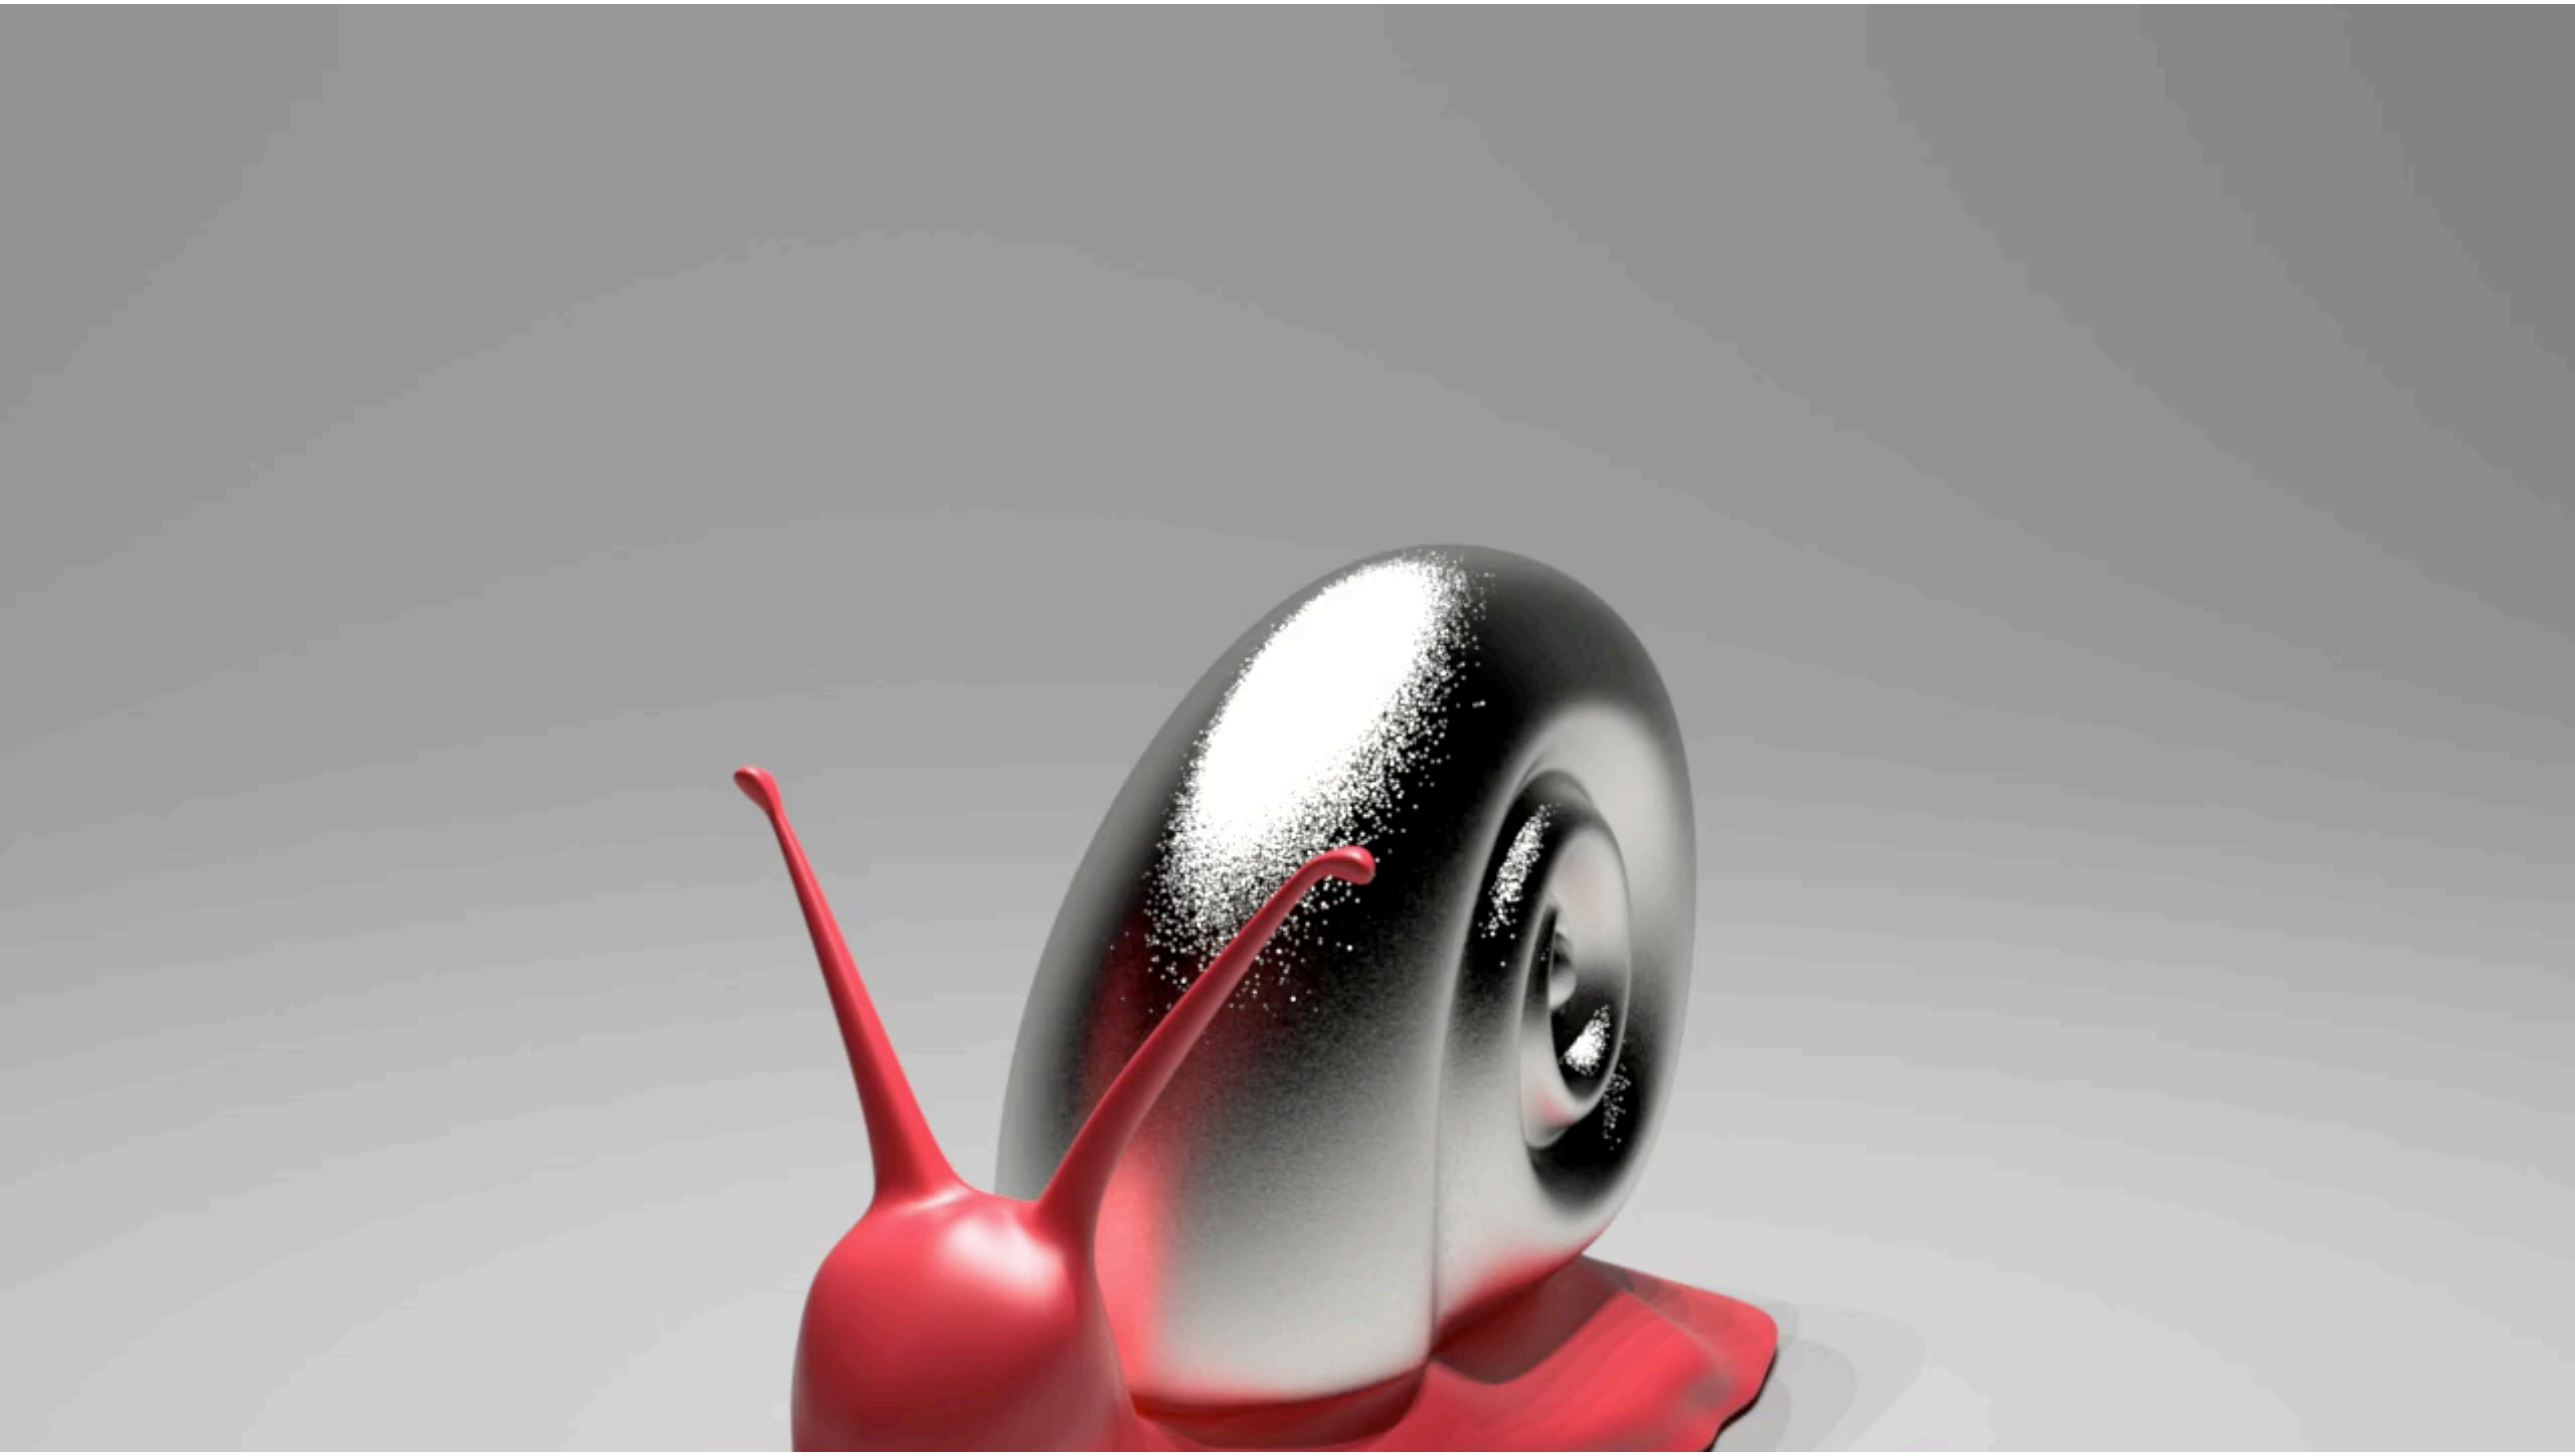

Different details

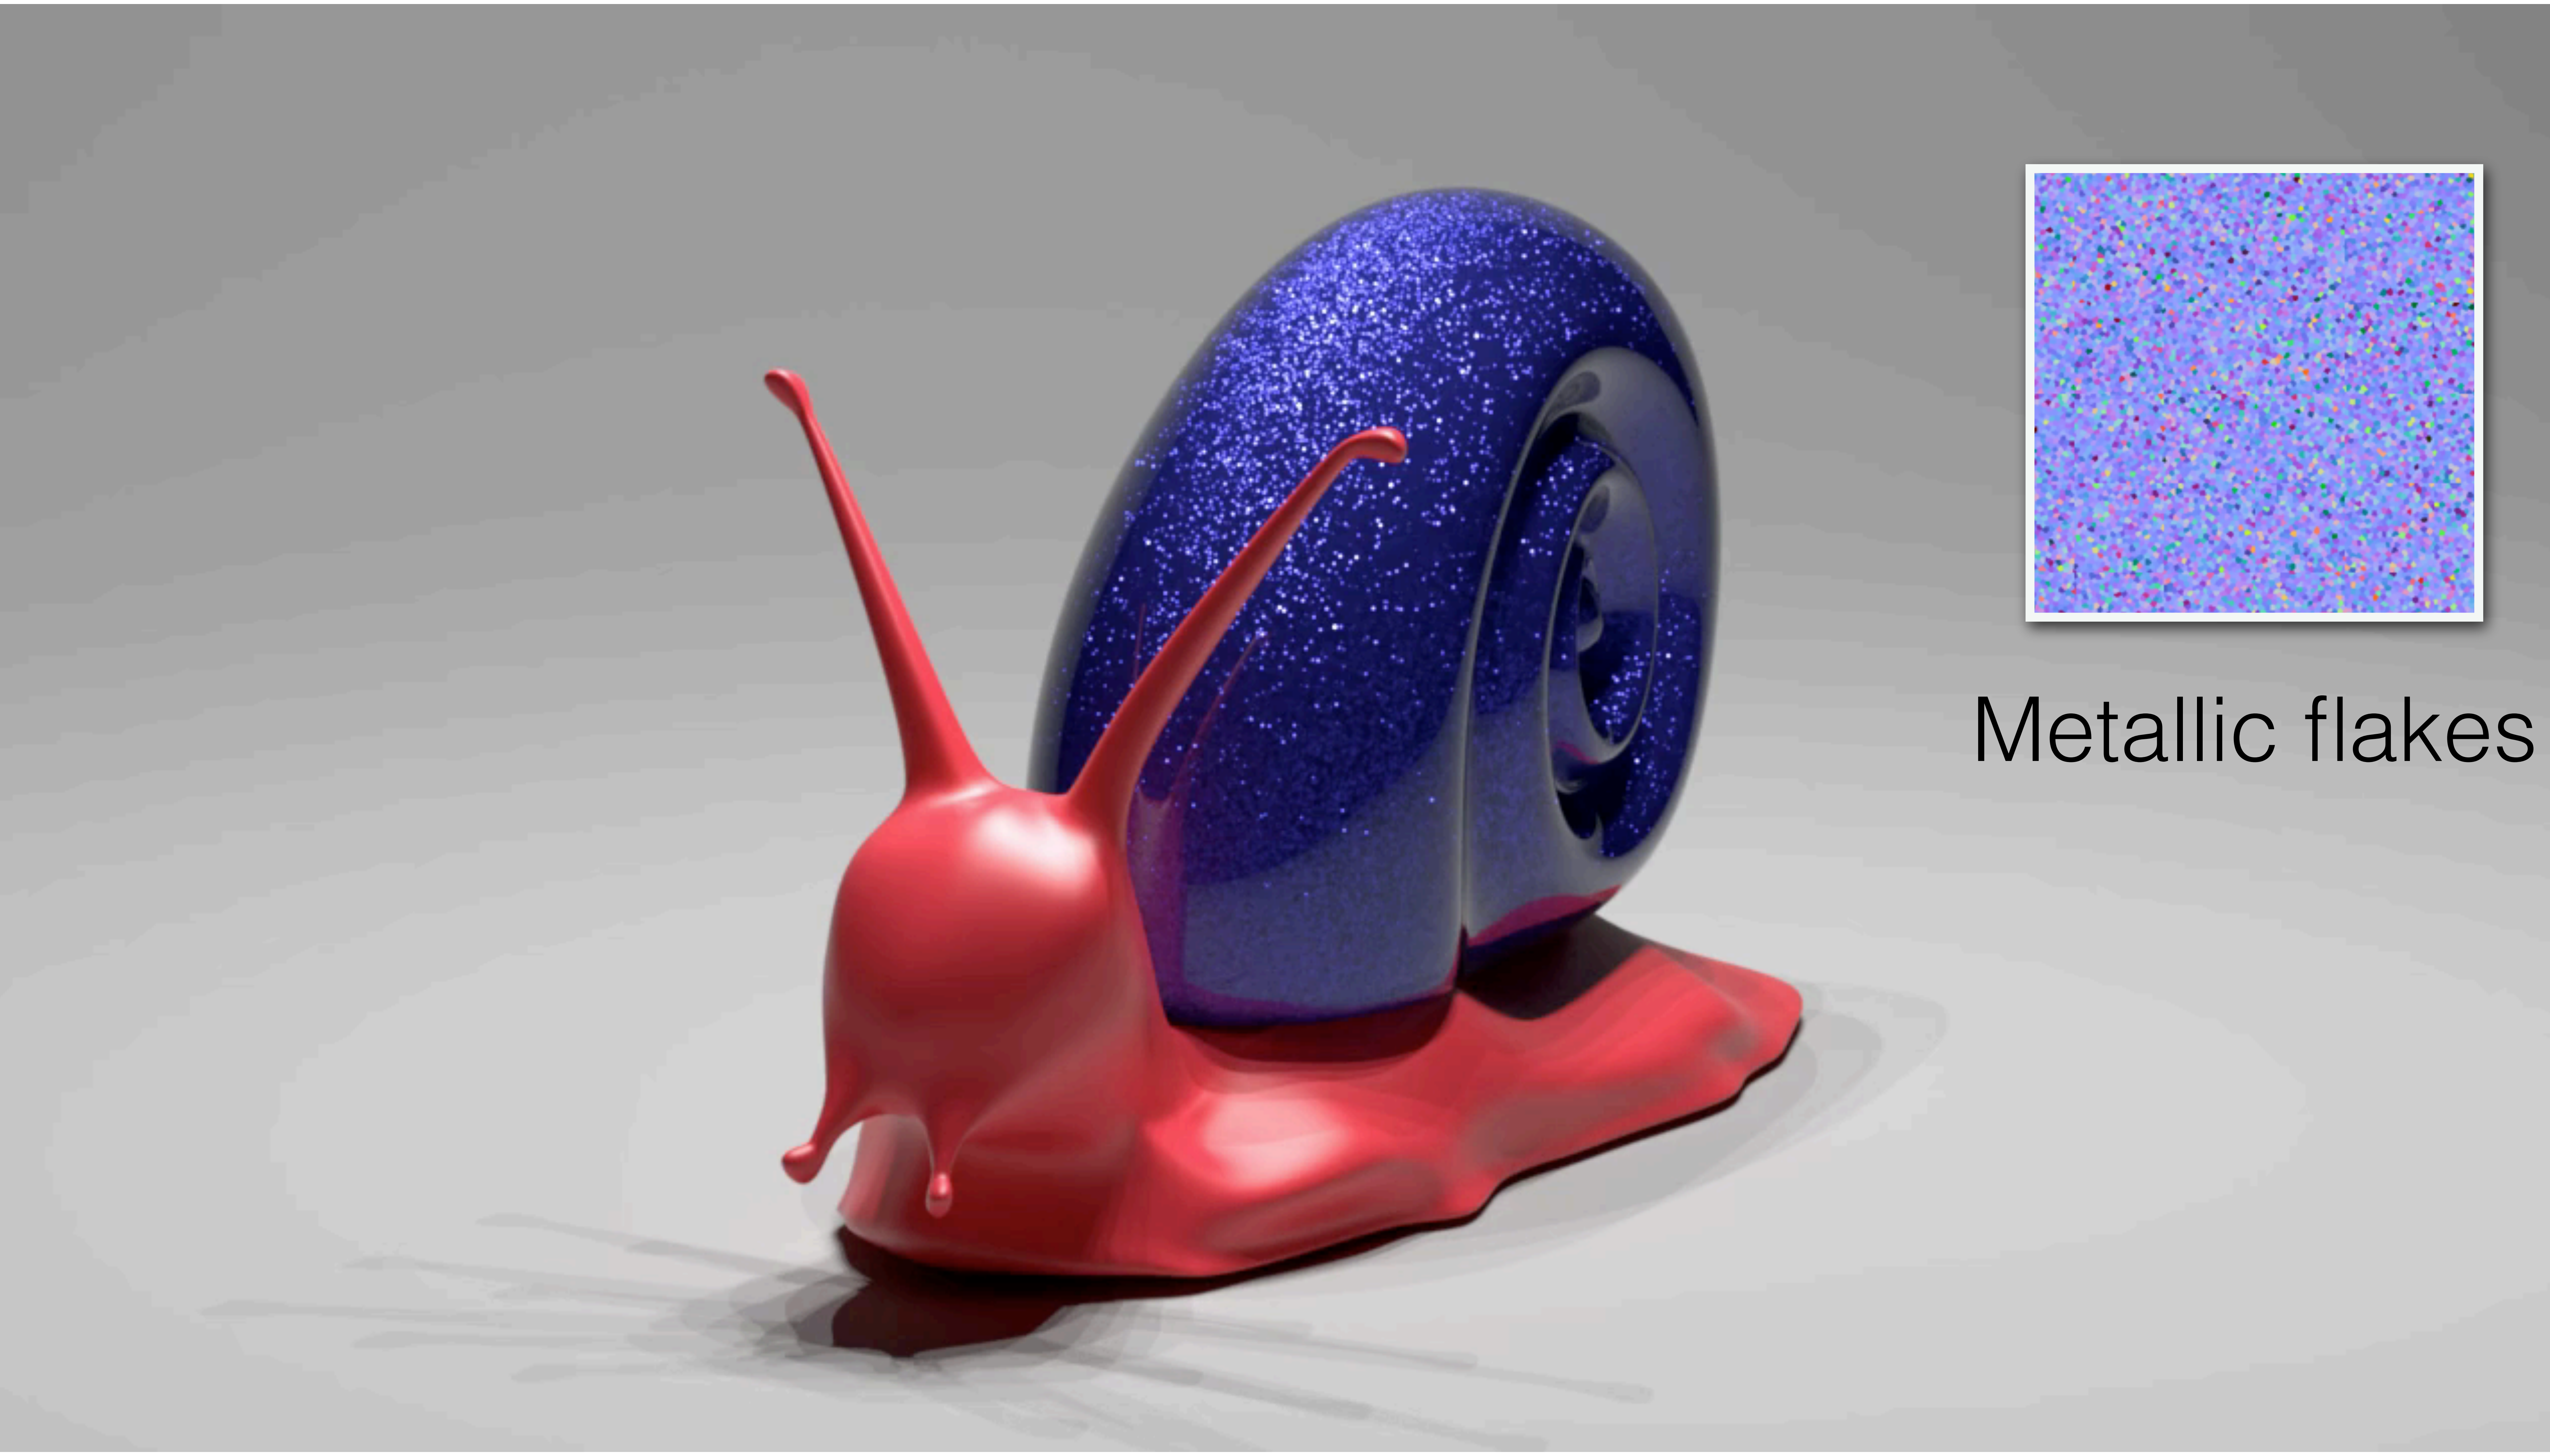

Metallic flakes

Rendering? Too difficult!

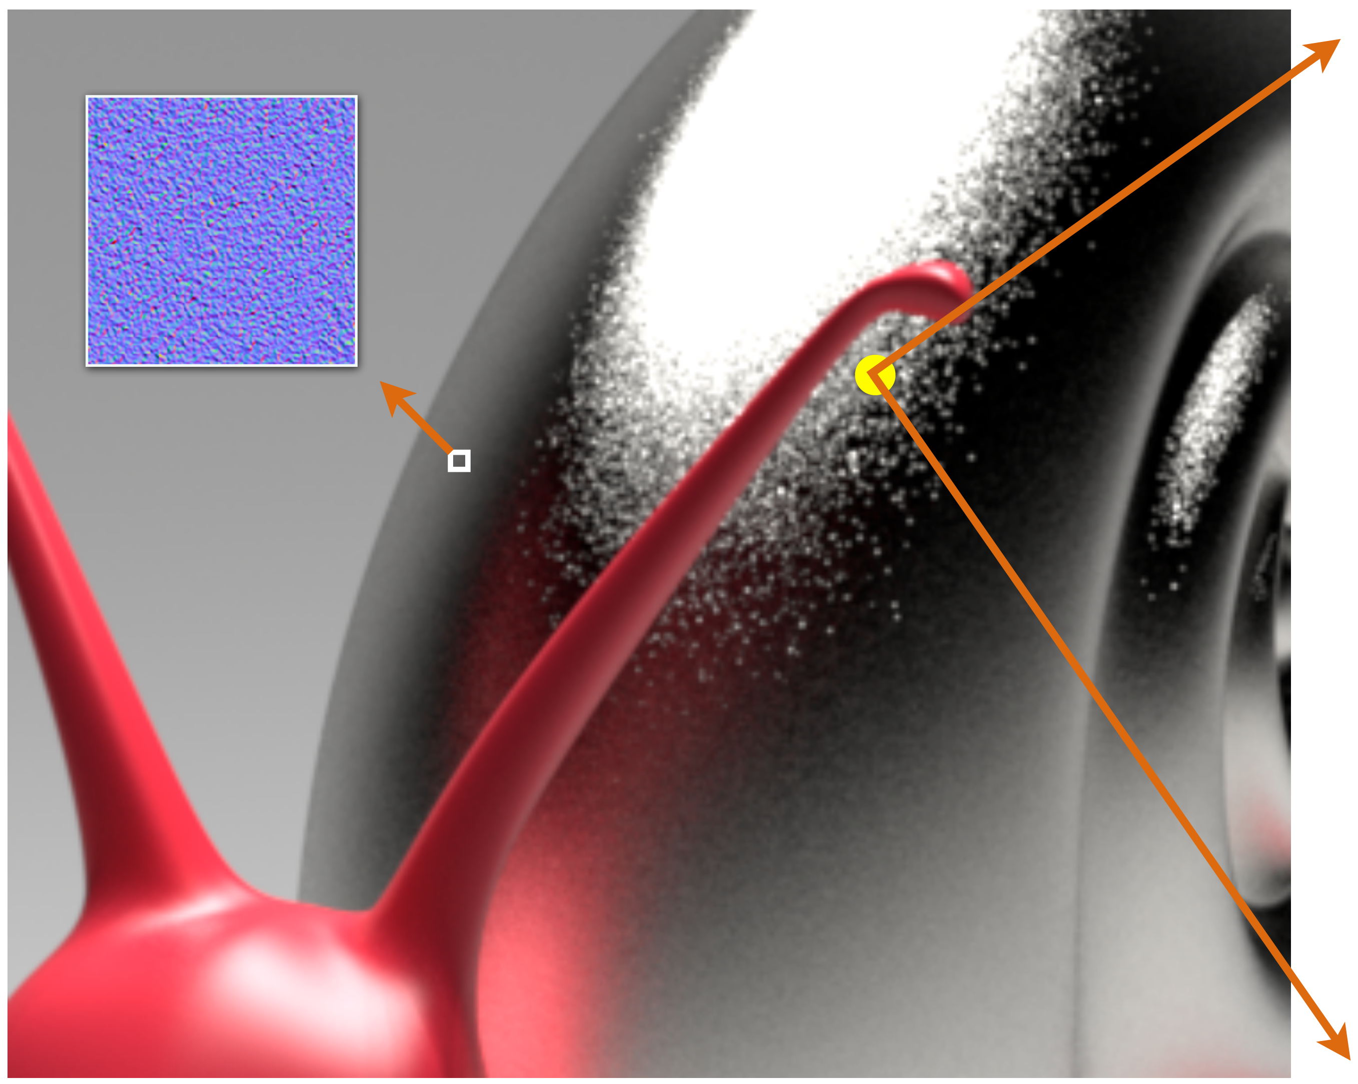

our result

zoom of
a single pixel

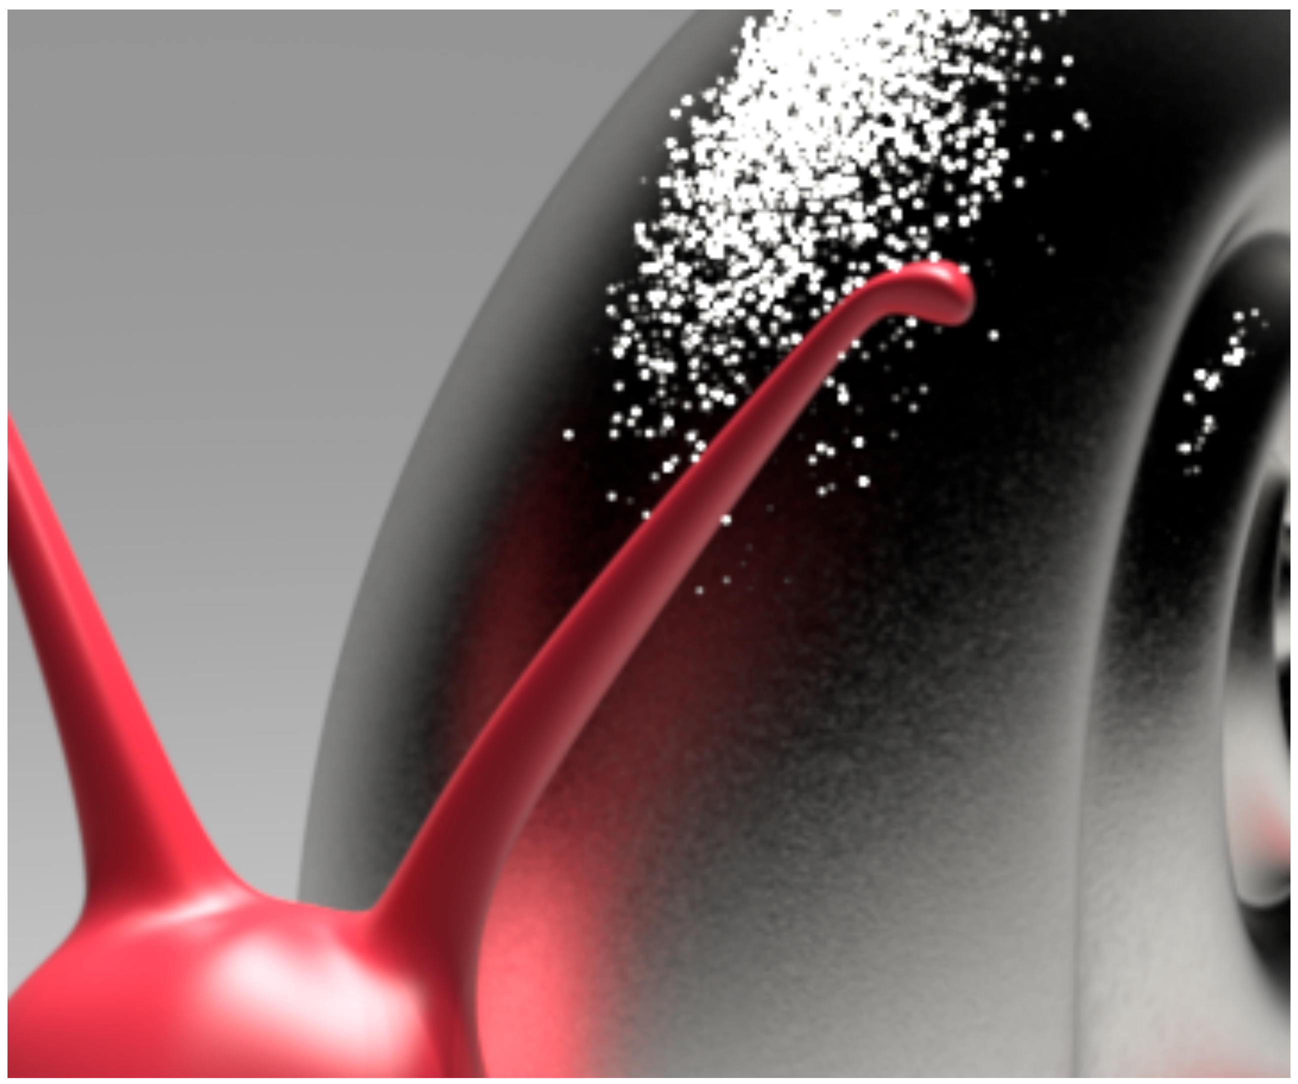

naive sampling (2h)
(≫ 21.3 days to converge)

Difficult path sampling problem

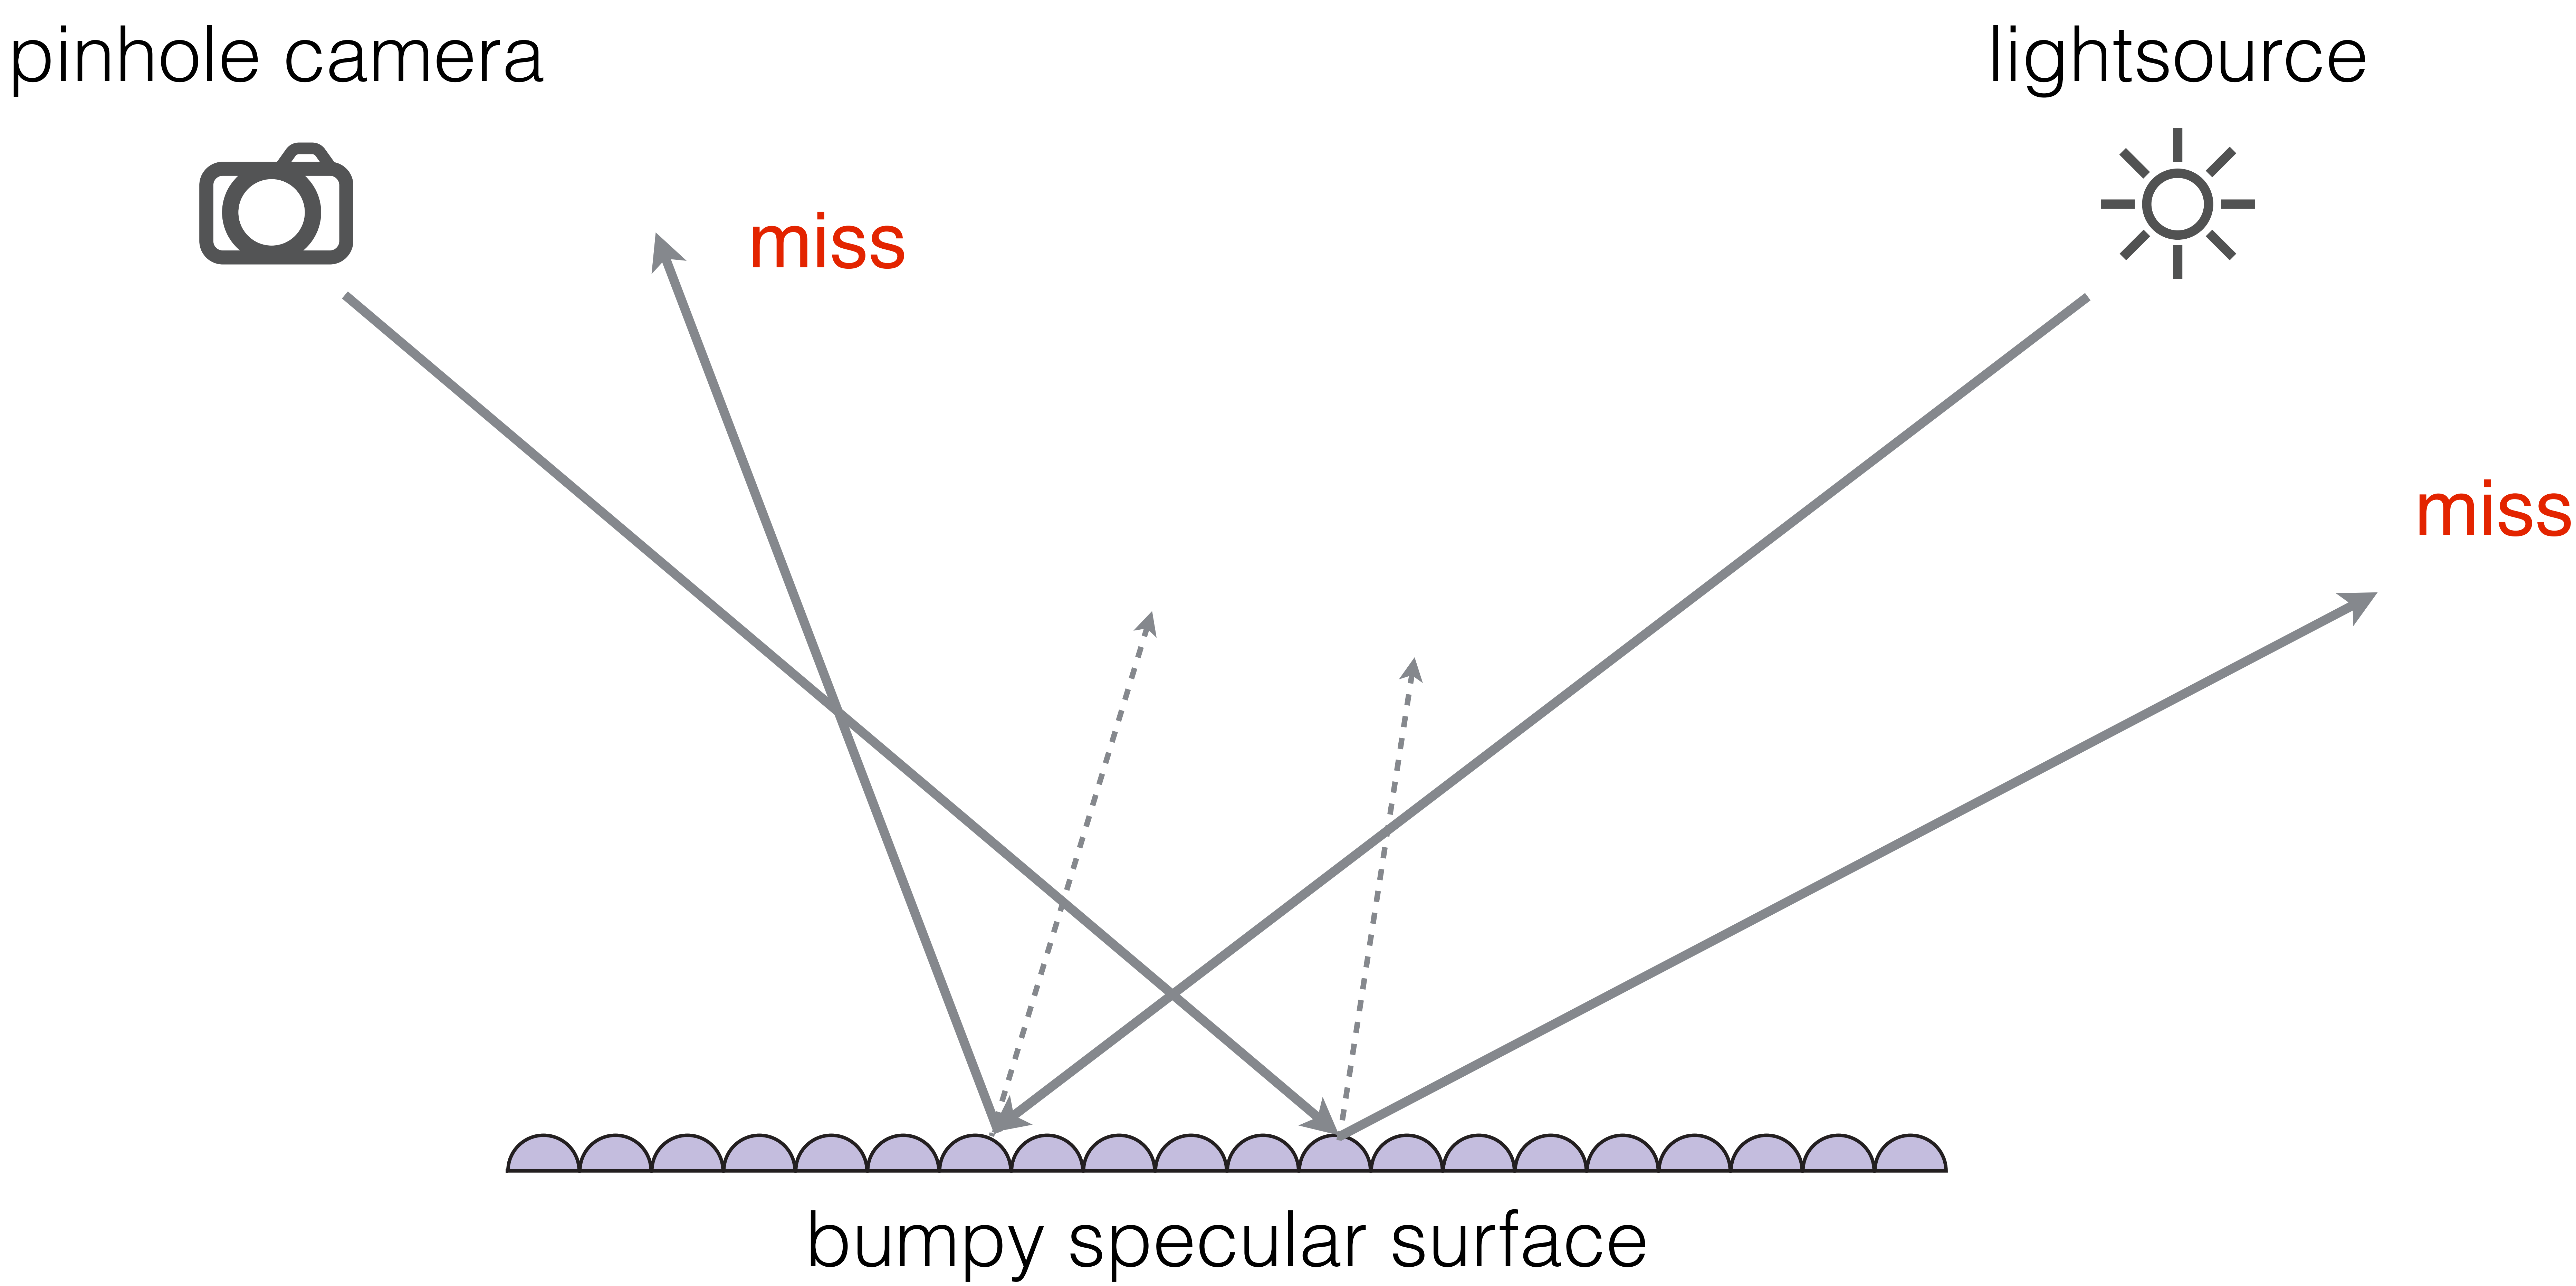

Solution: BRDF over a pixel

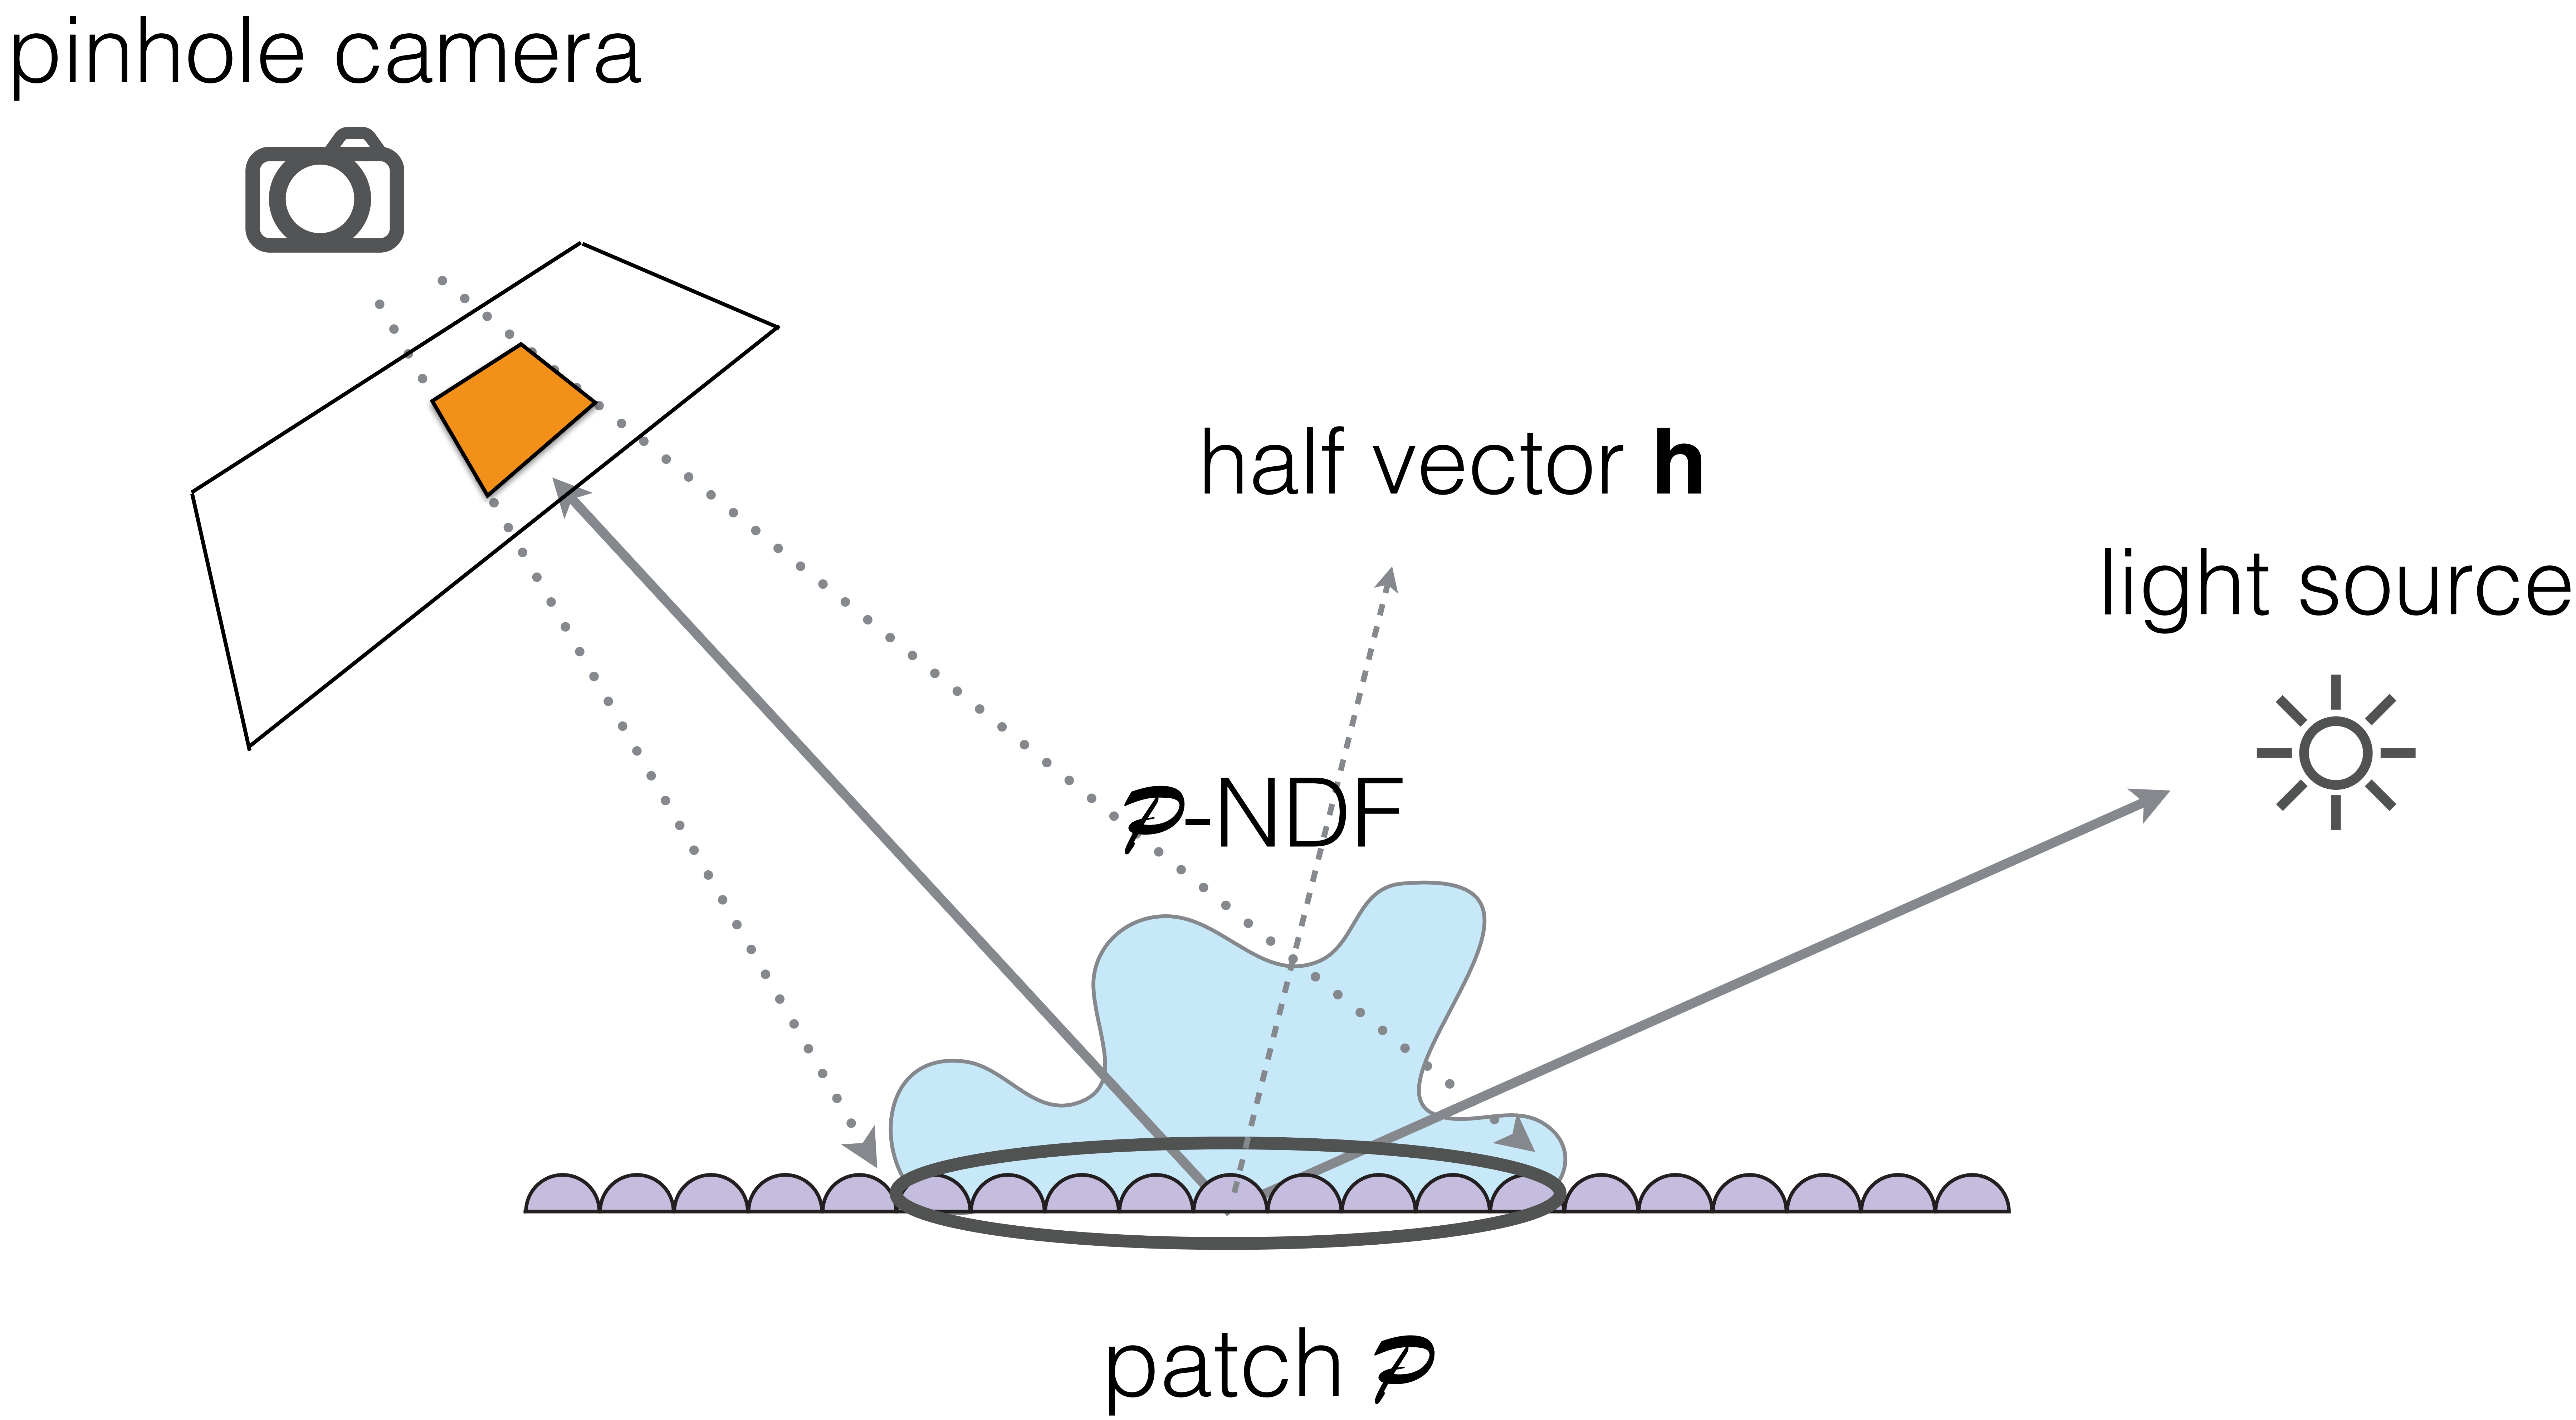

p-NDFs have sharp features

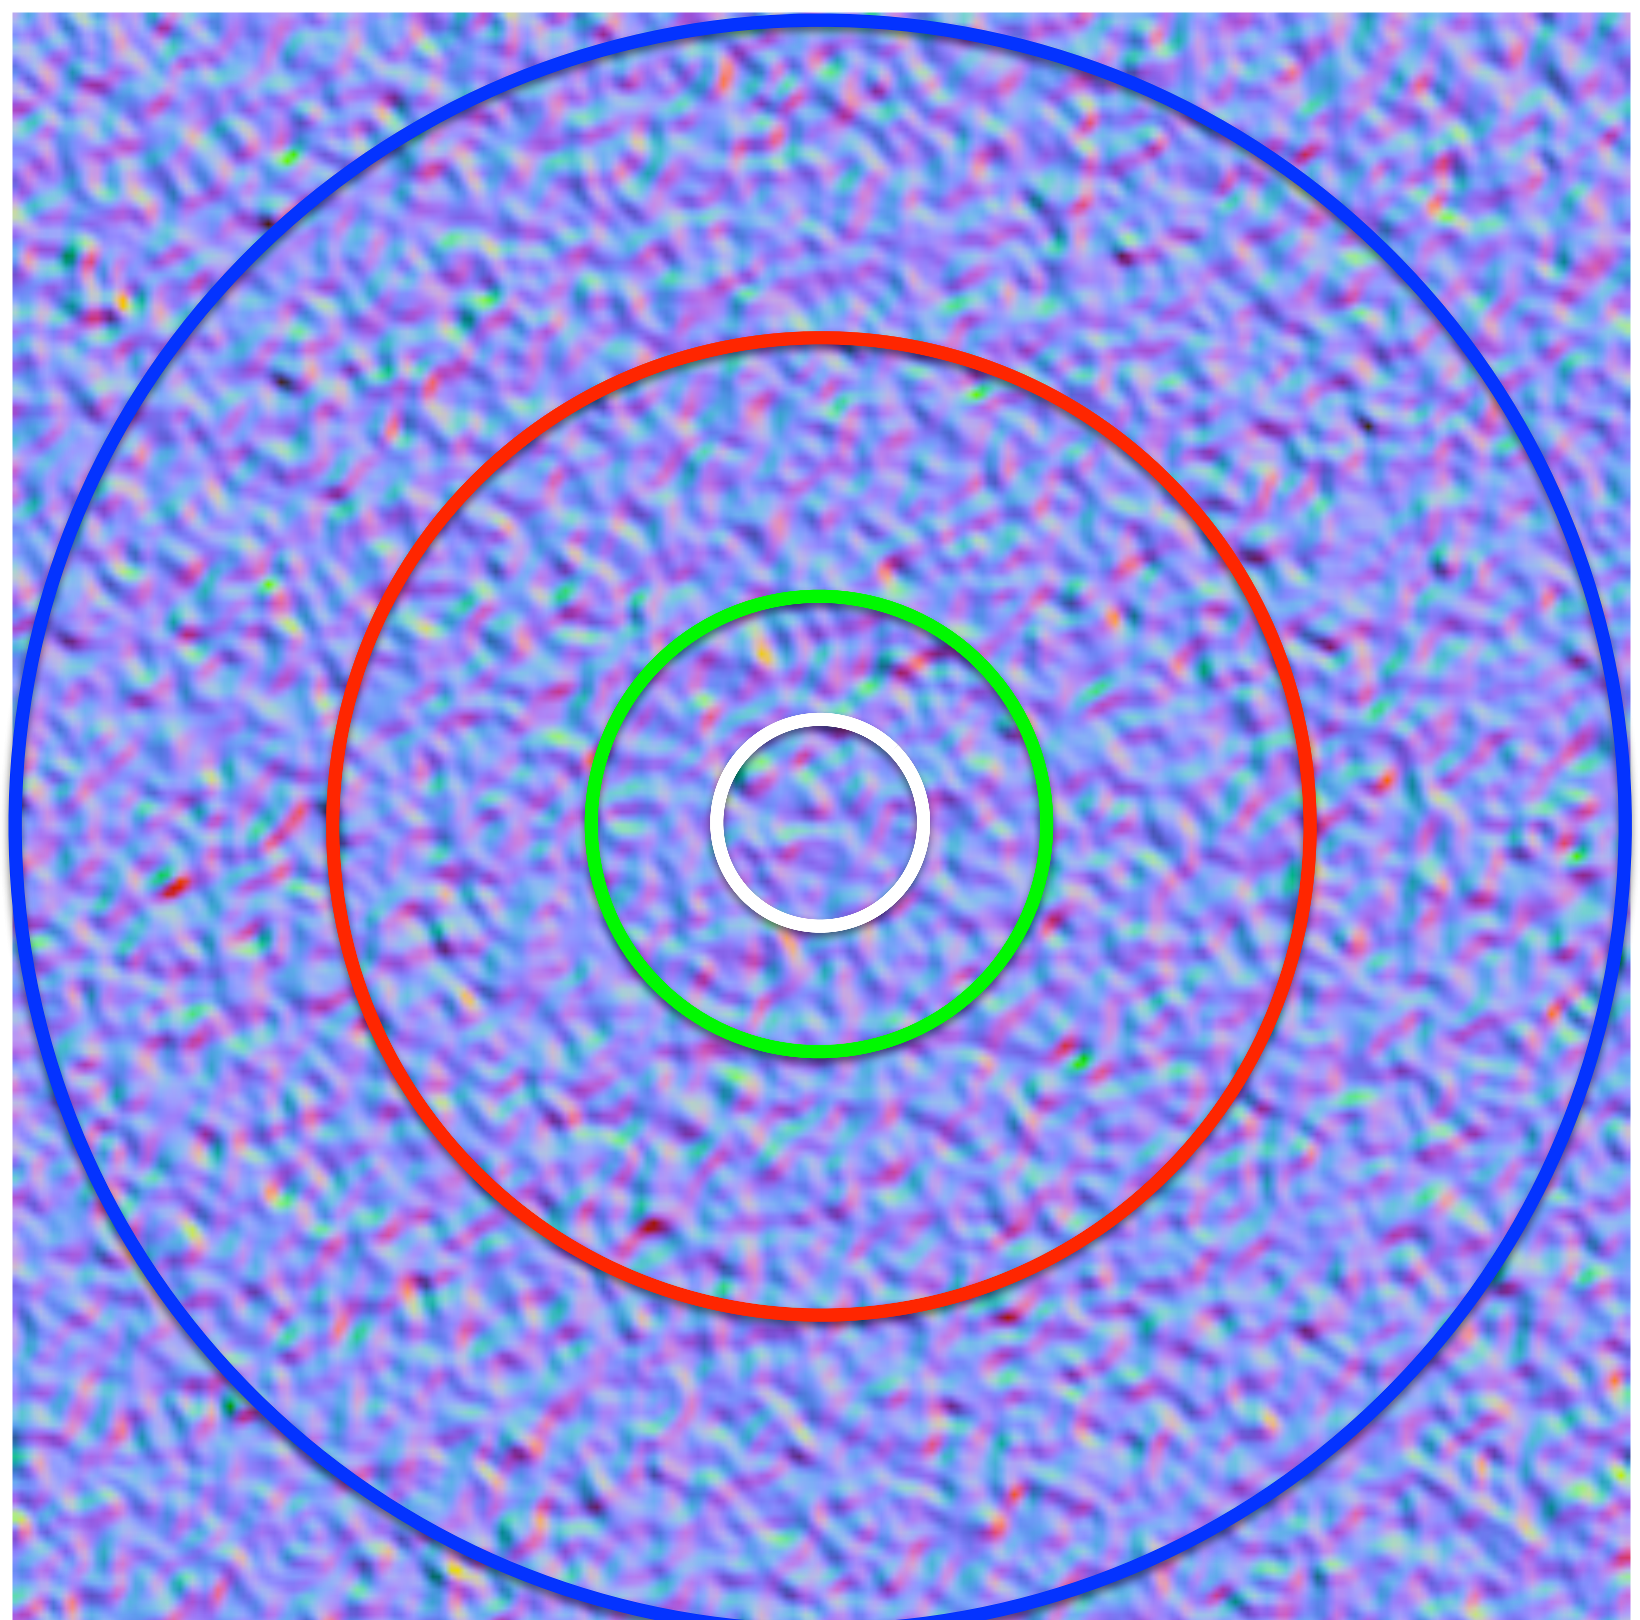

normal map

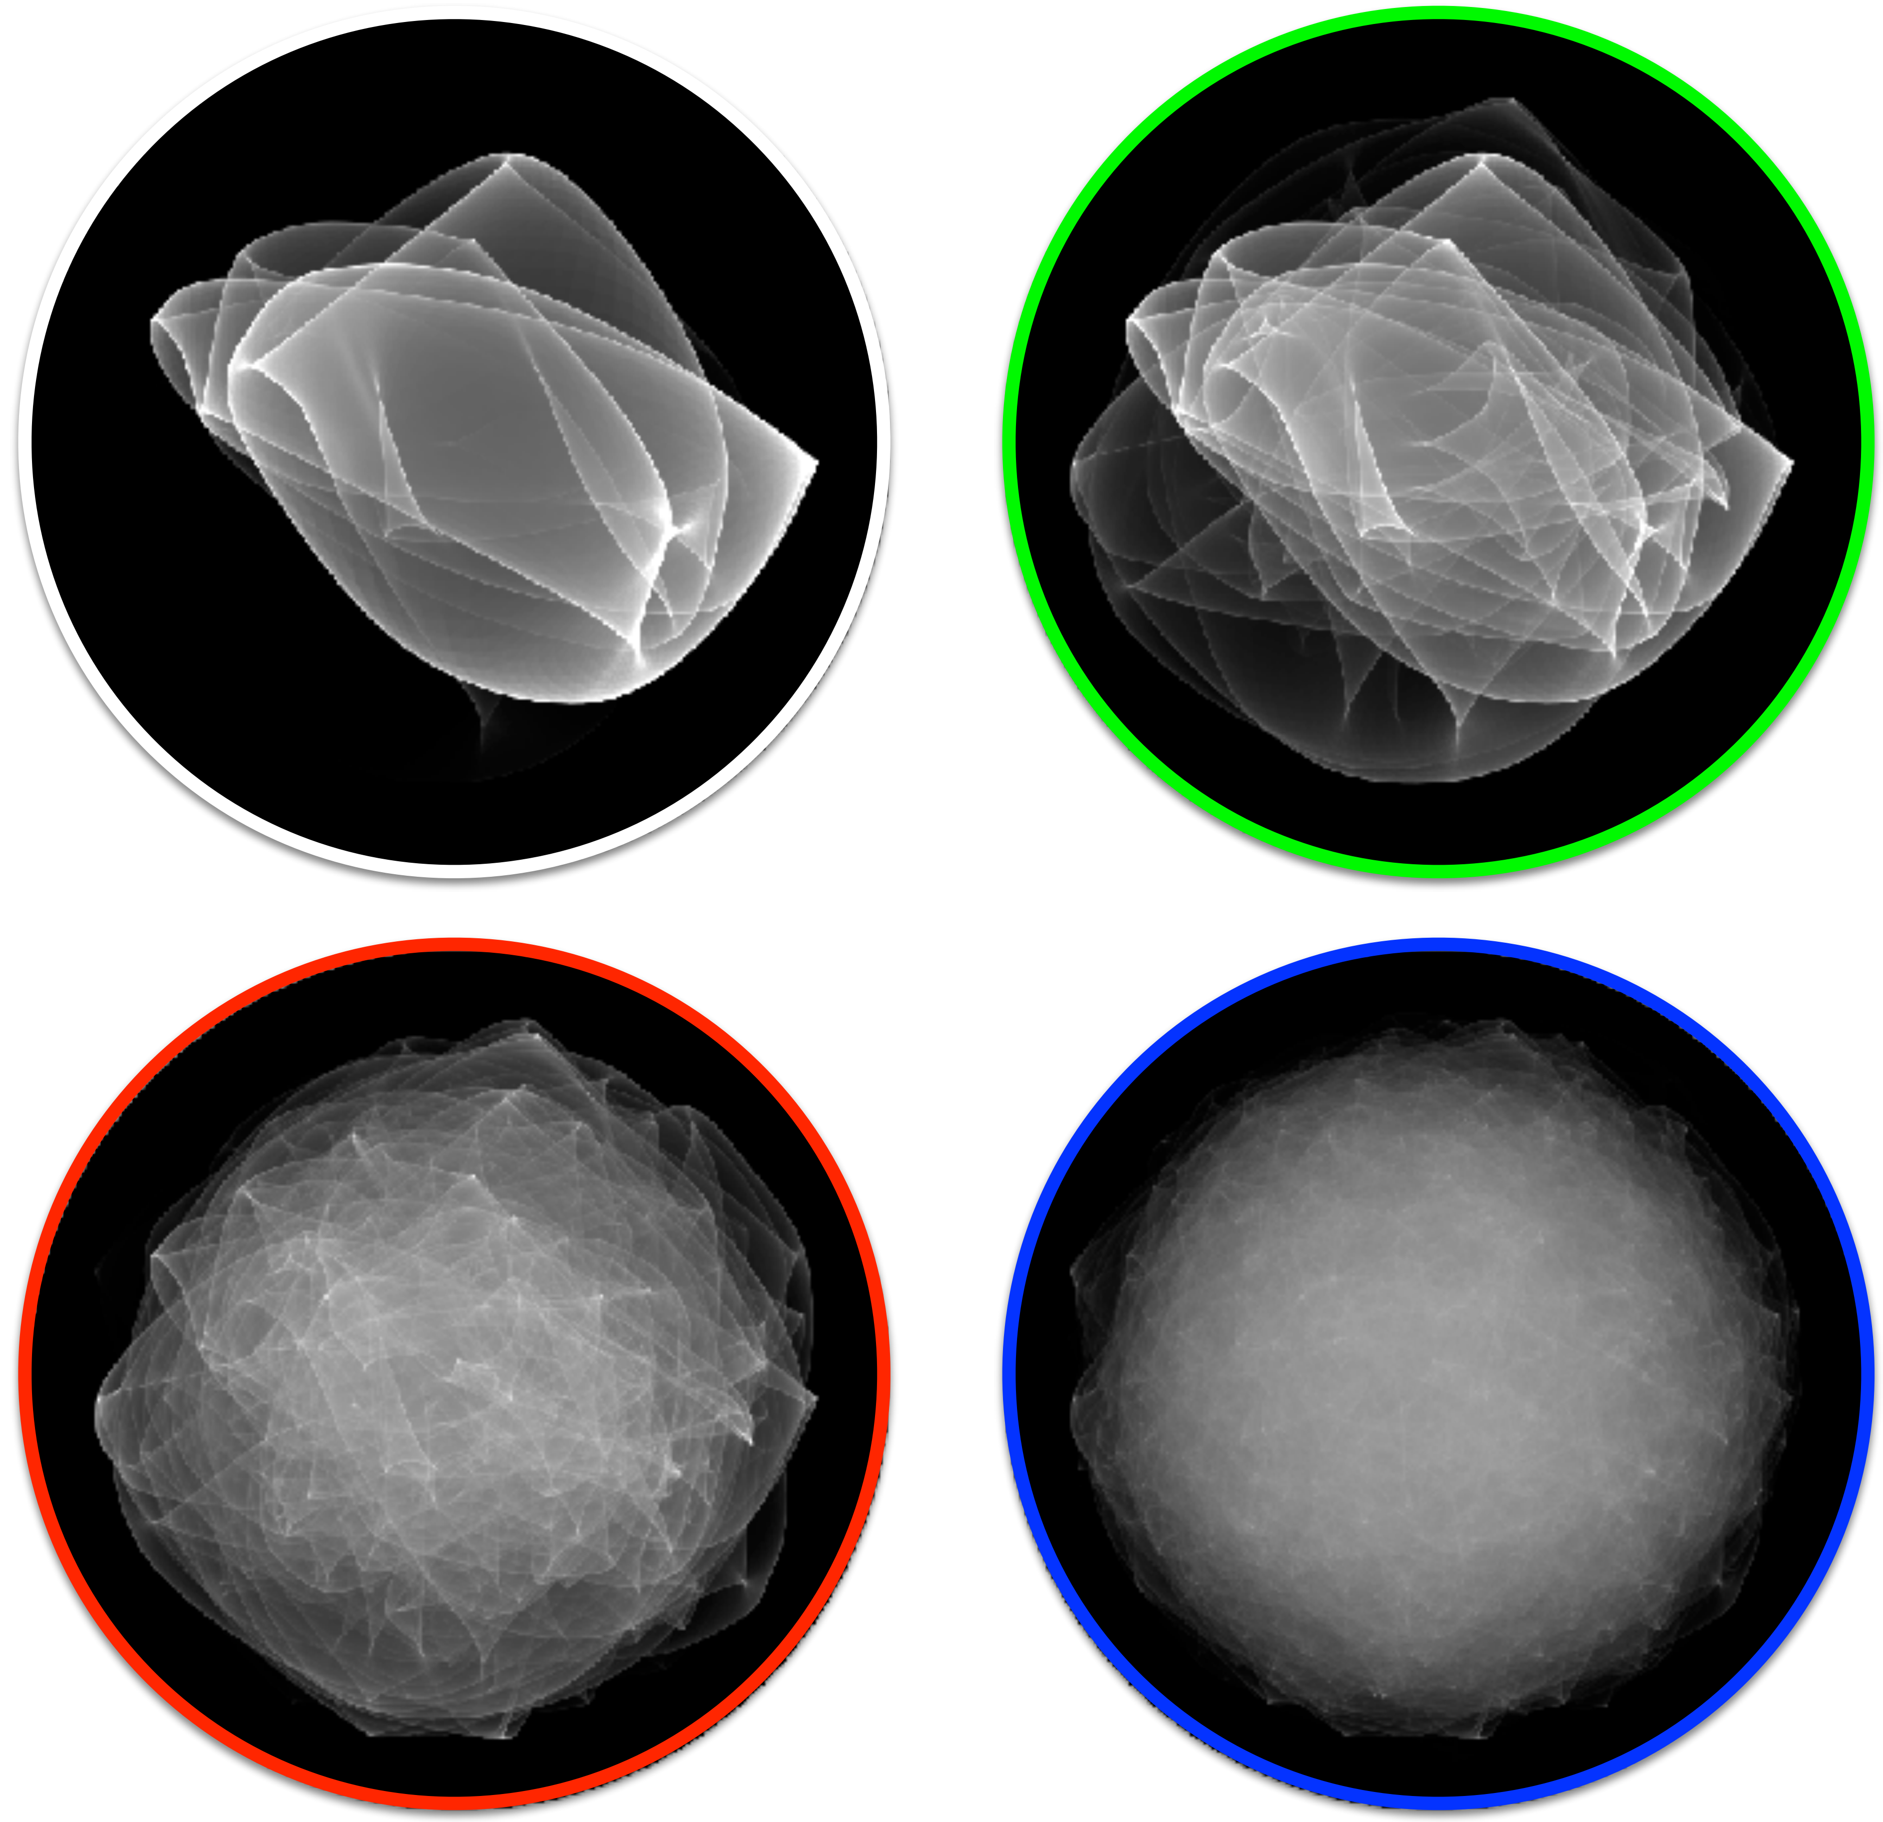

$\mathcal{P}$-NDFs

p-NDF shapes

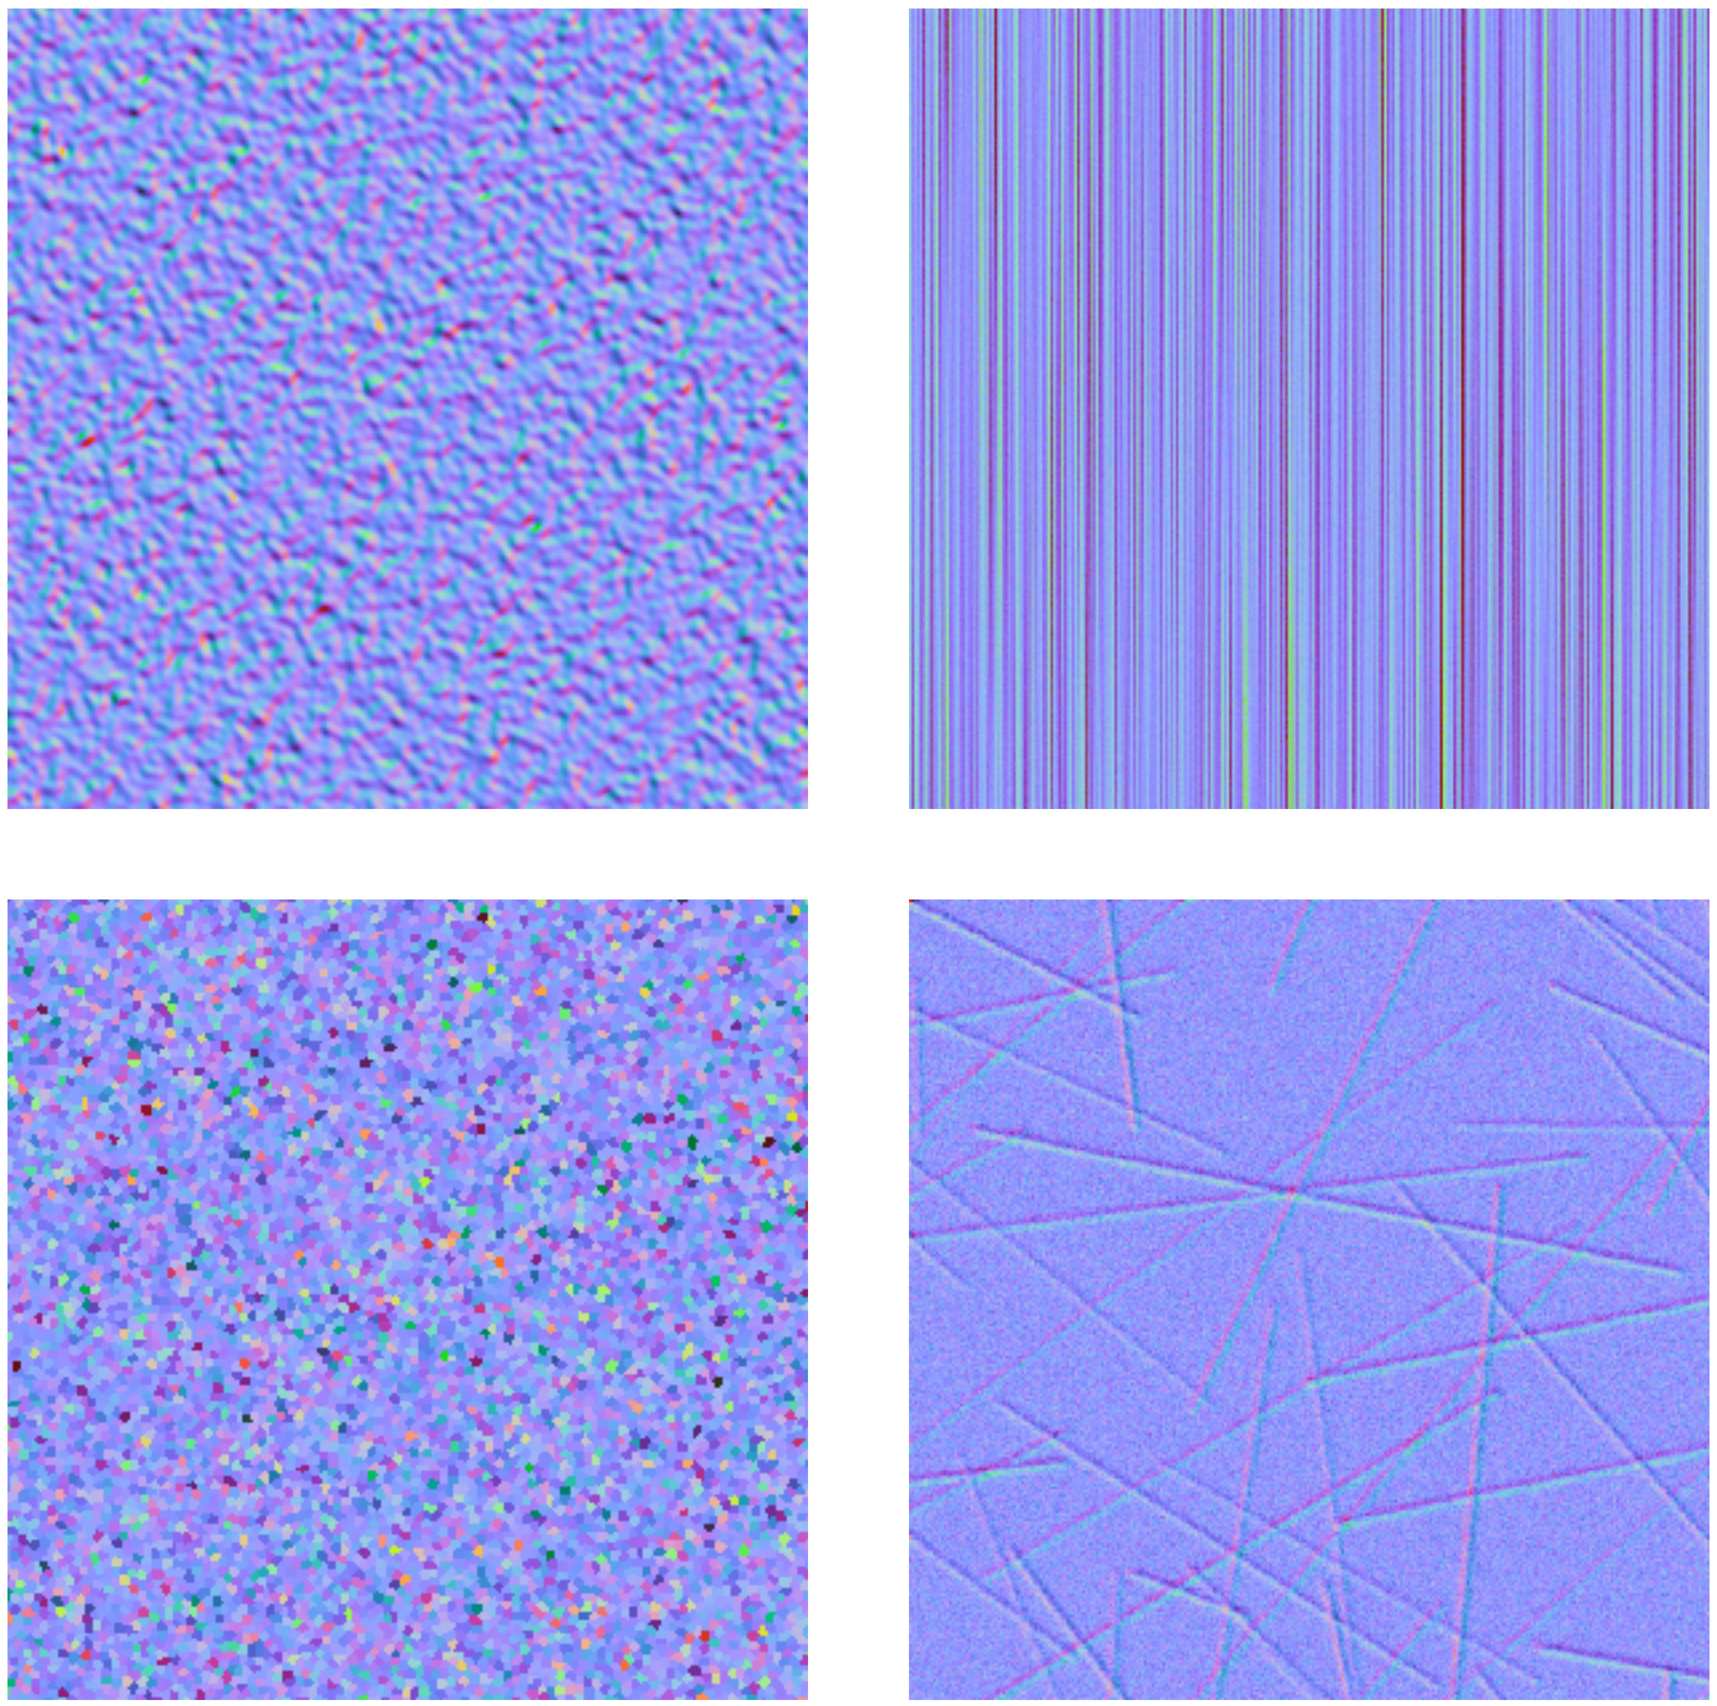

normal maps

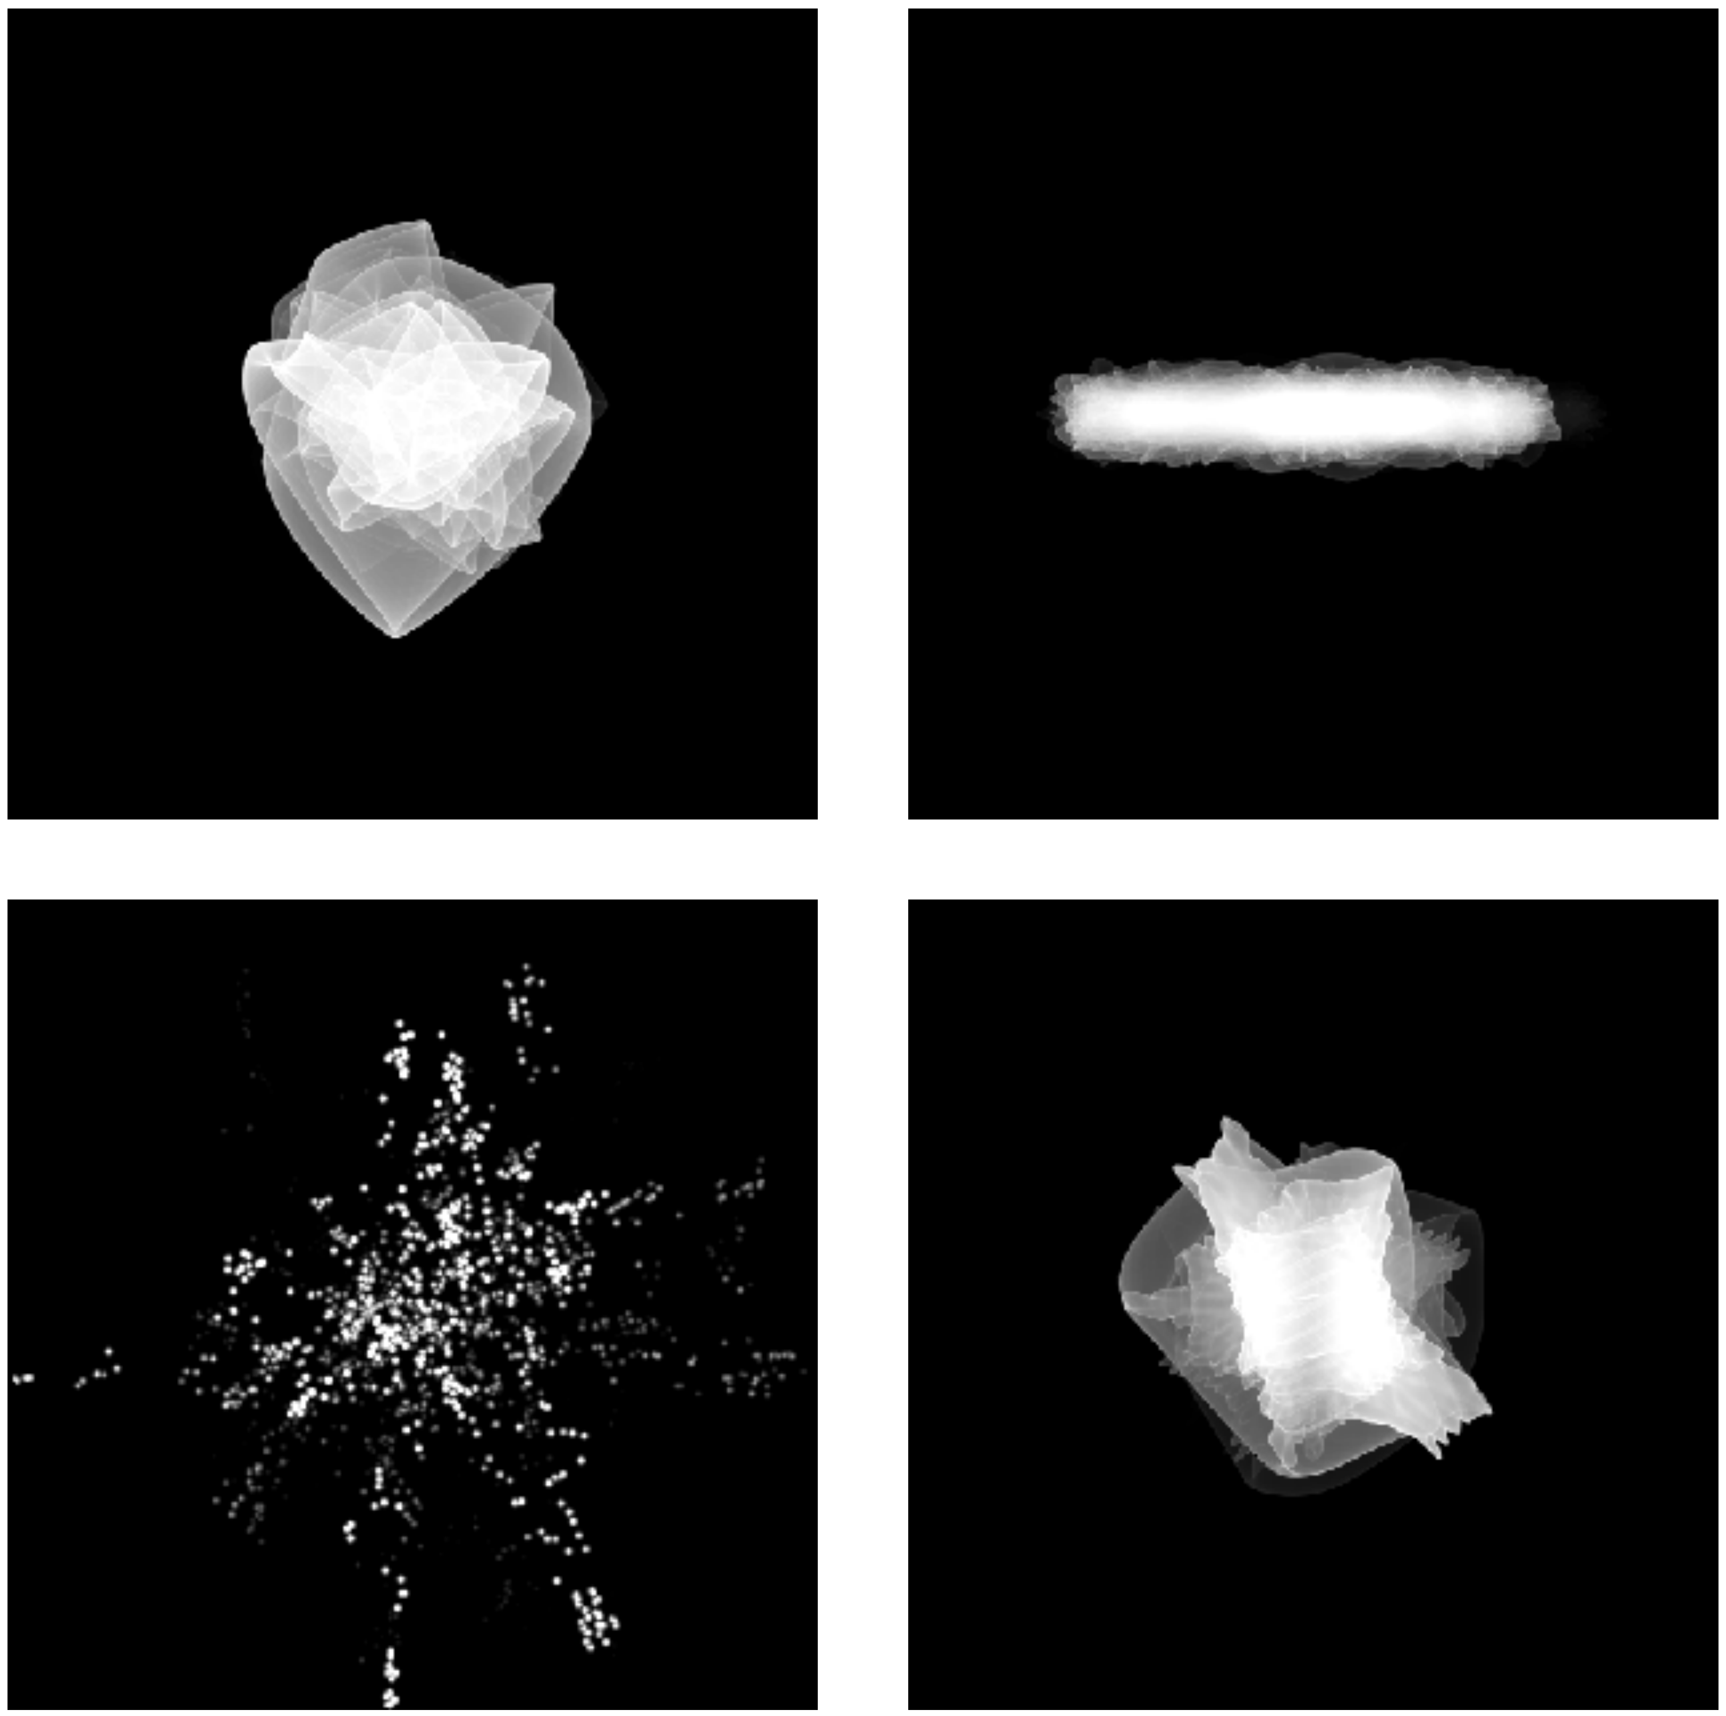

$\mathcal{P}$-NDFs

Blender

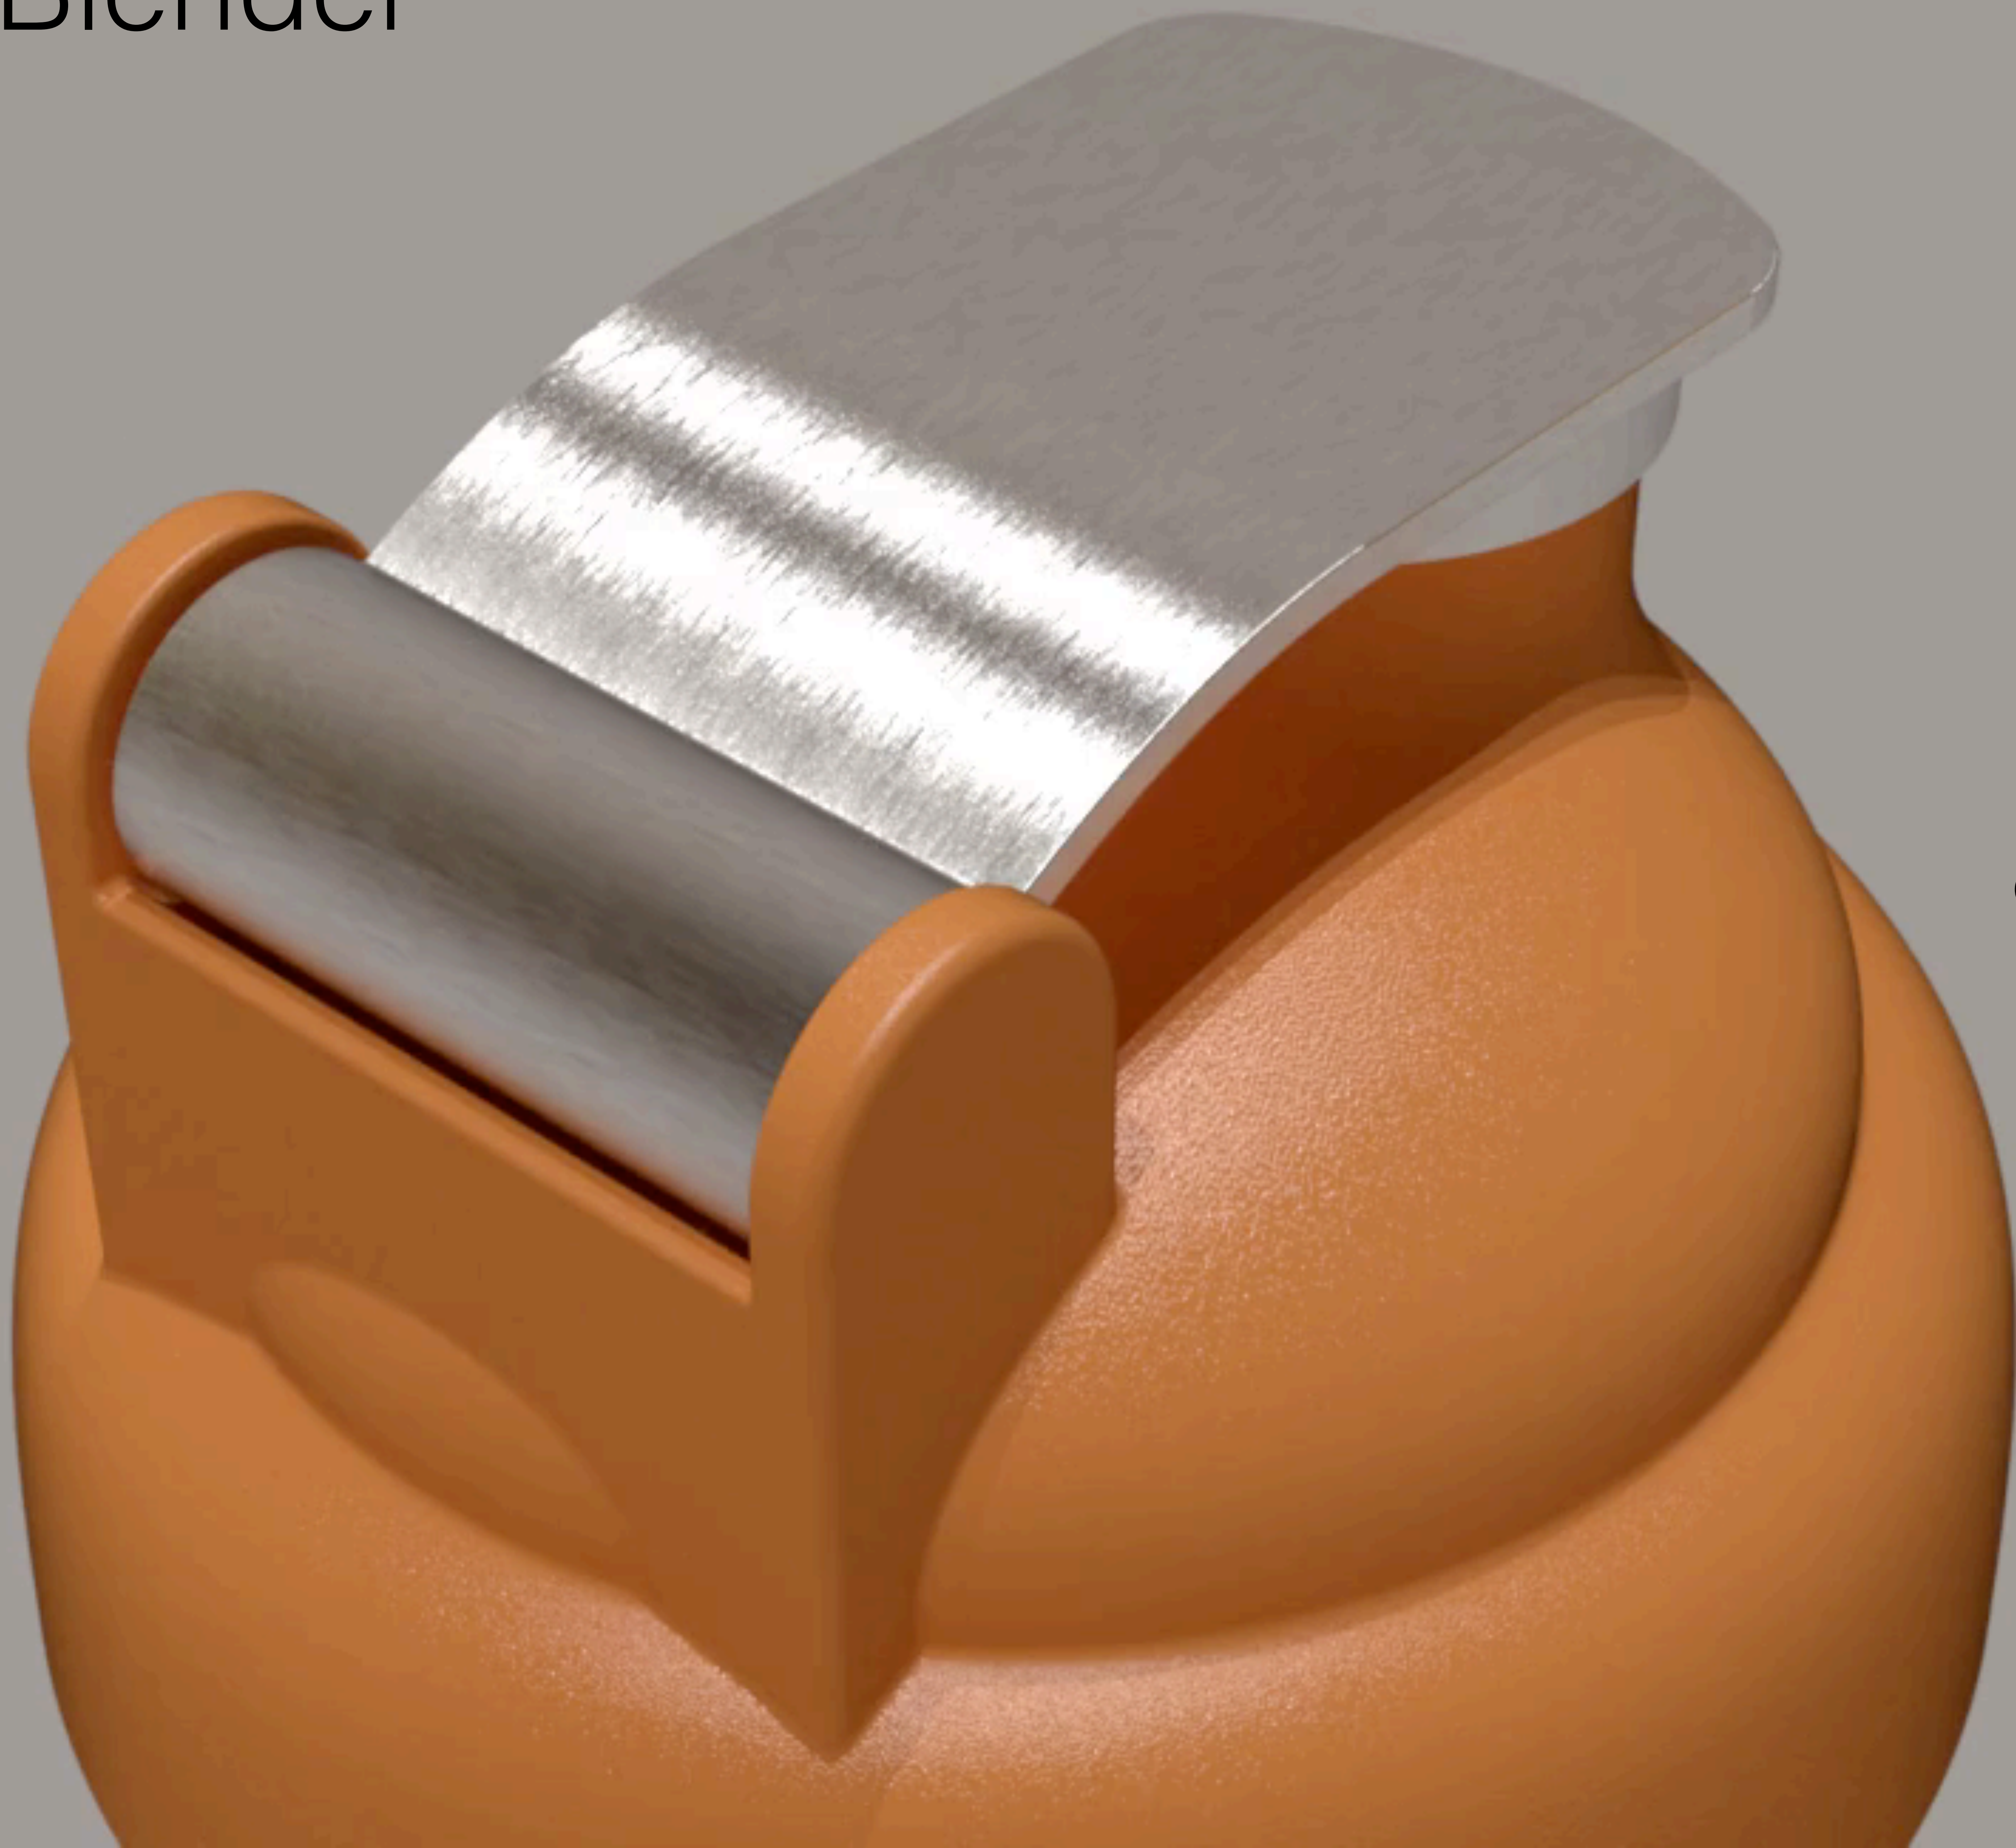

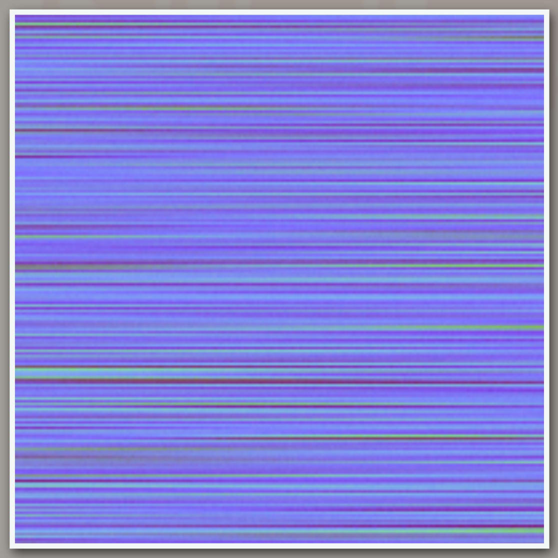

brushed metal

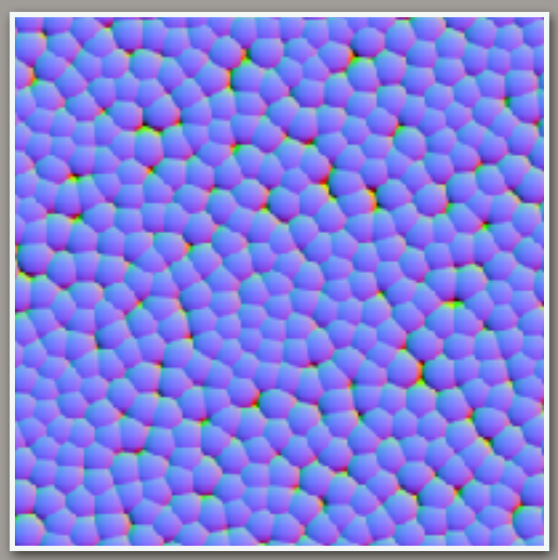

ellipsoid bumps

Blender: Zoom

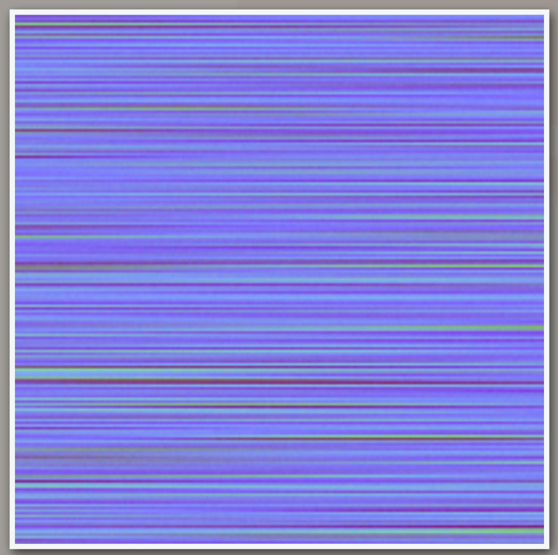

brushed metal

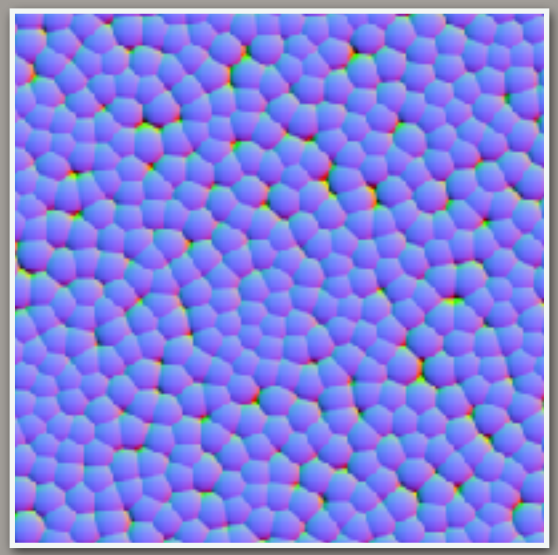

ellipsoid bumps

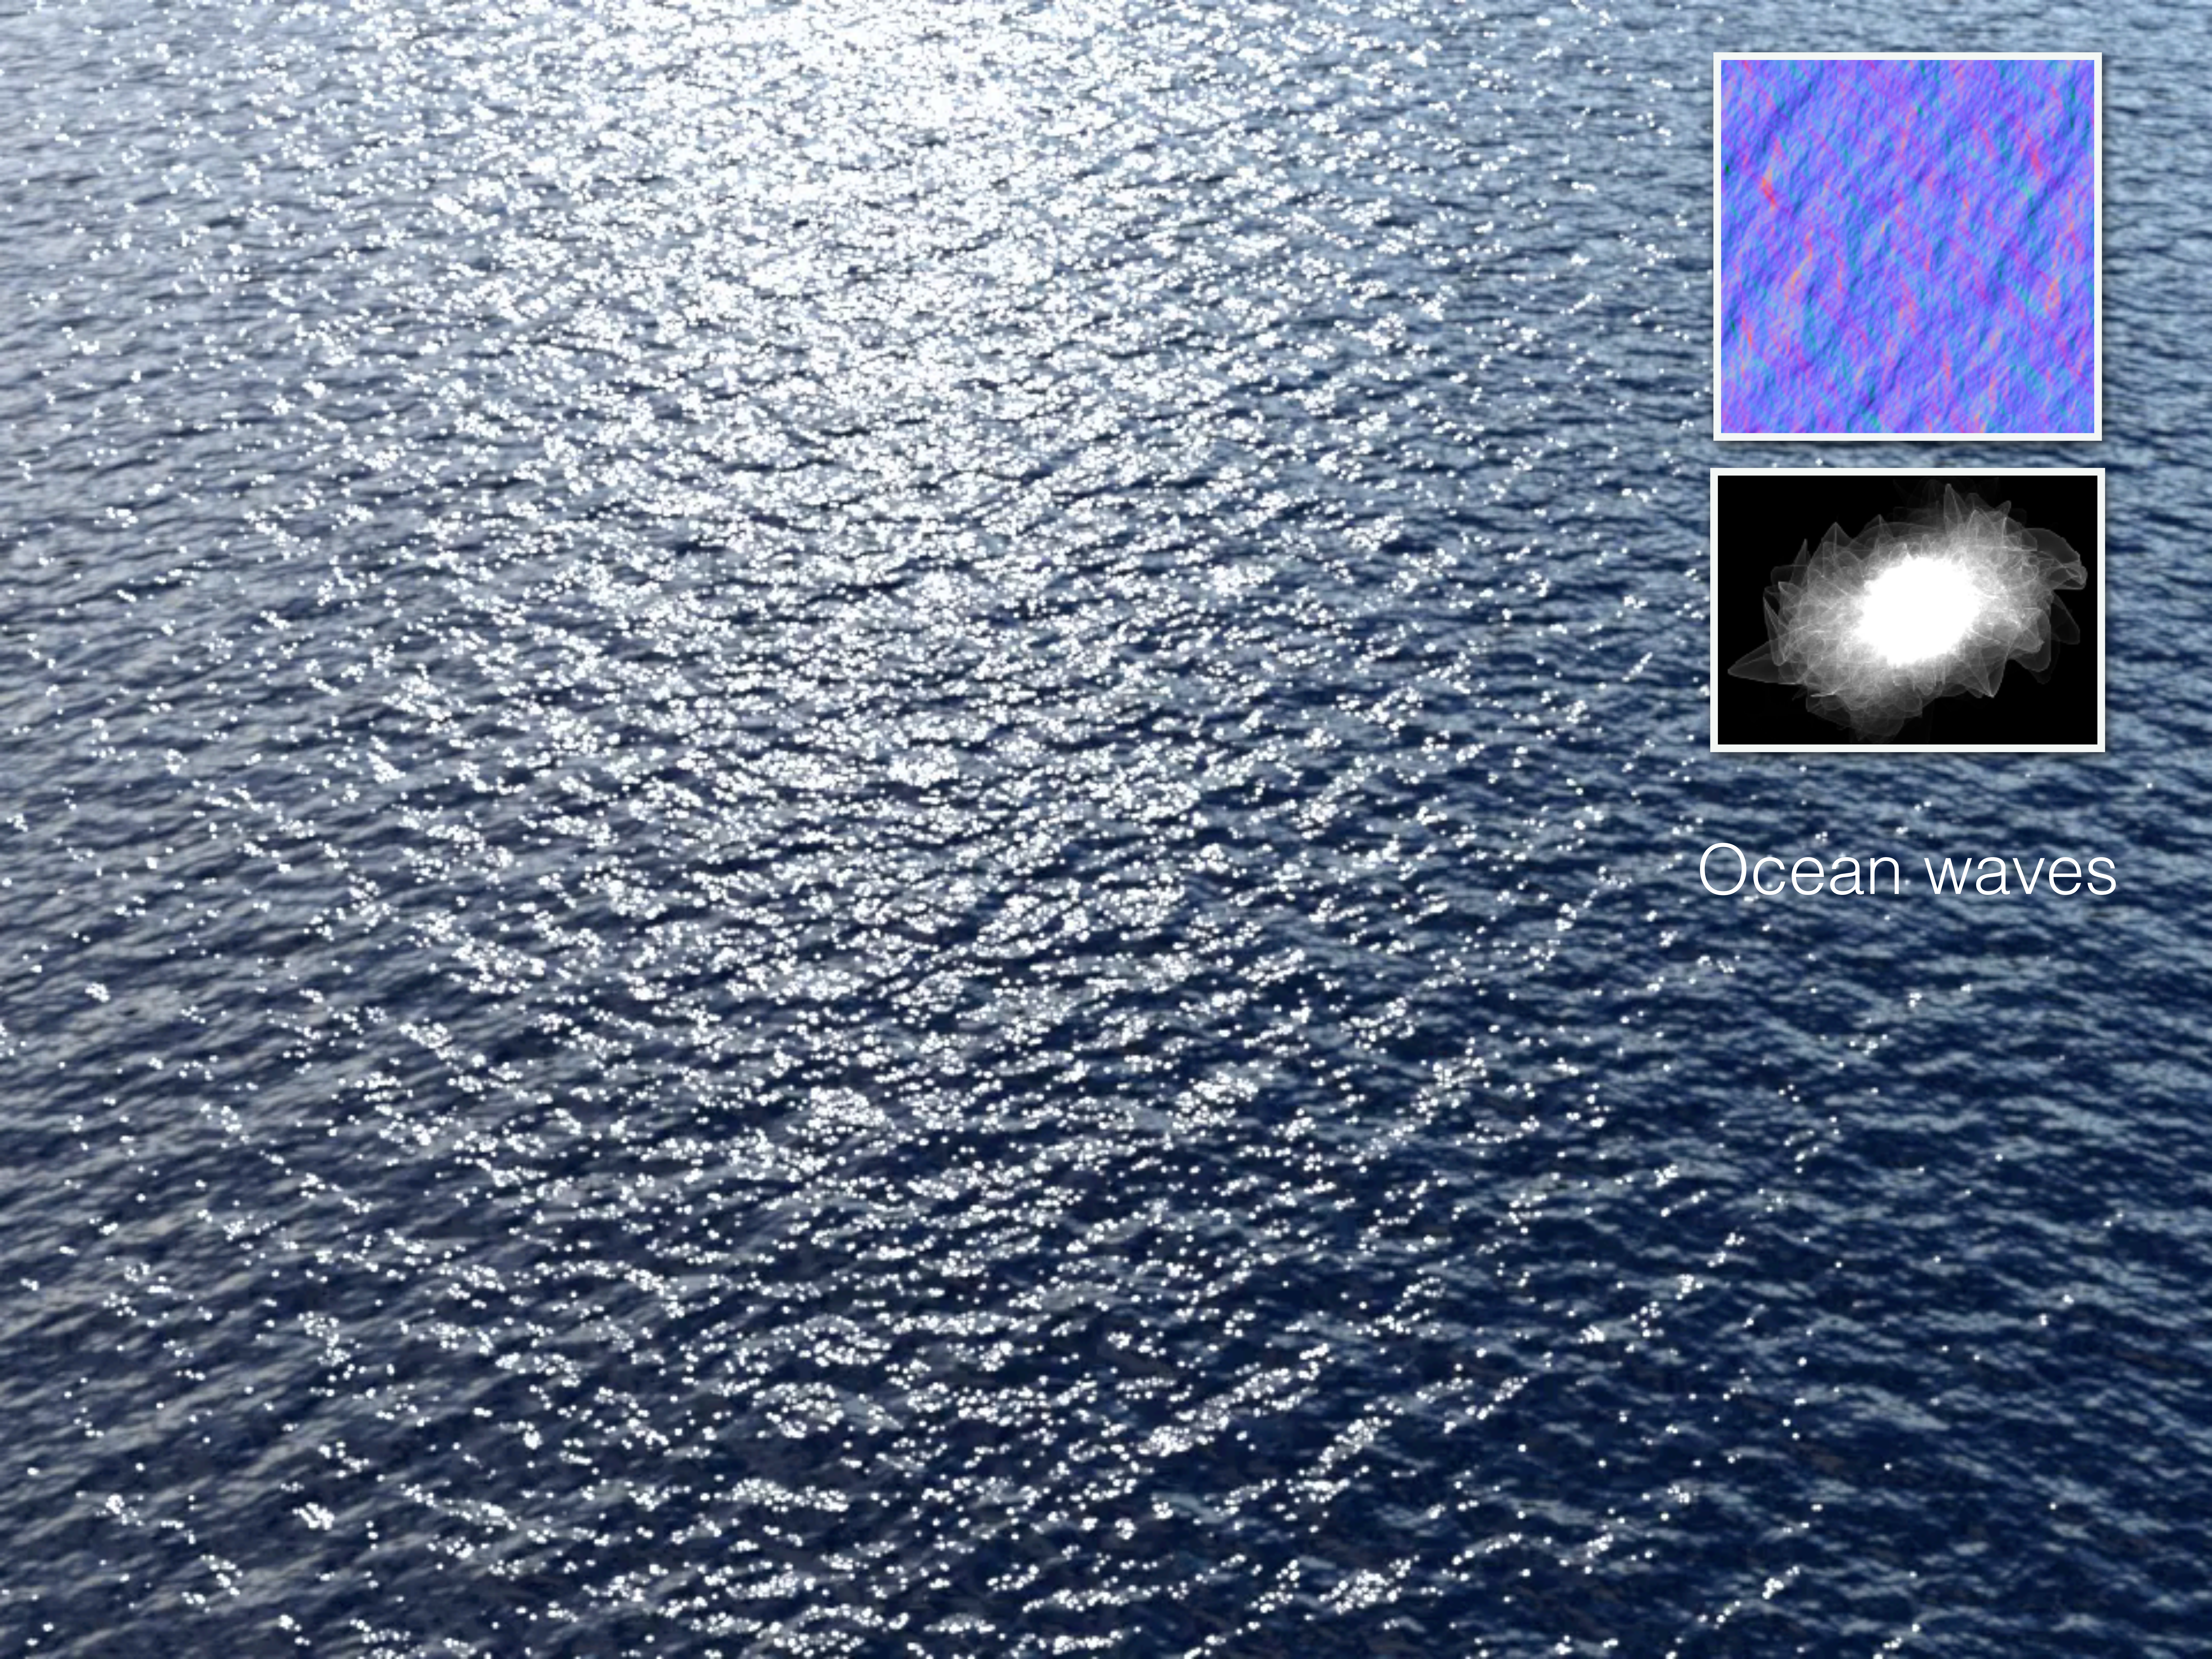

Ocean waves

Detailed / Glinty Material: Application

[Rise of the Tomb Raider. 2016 Square Enix]

Recent Trend: Wave Optics

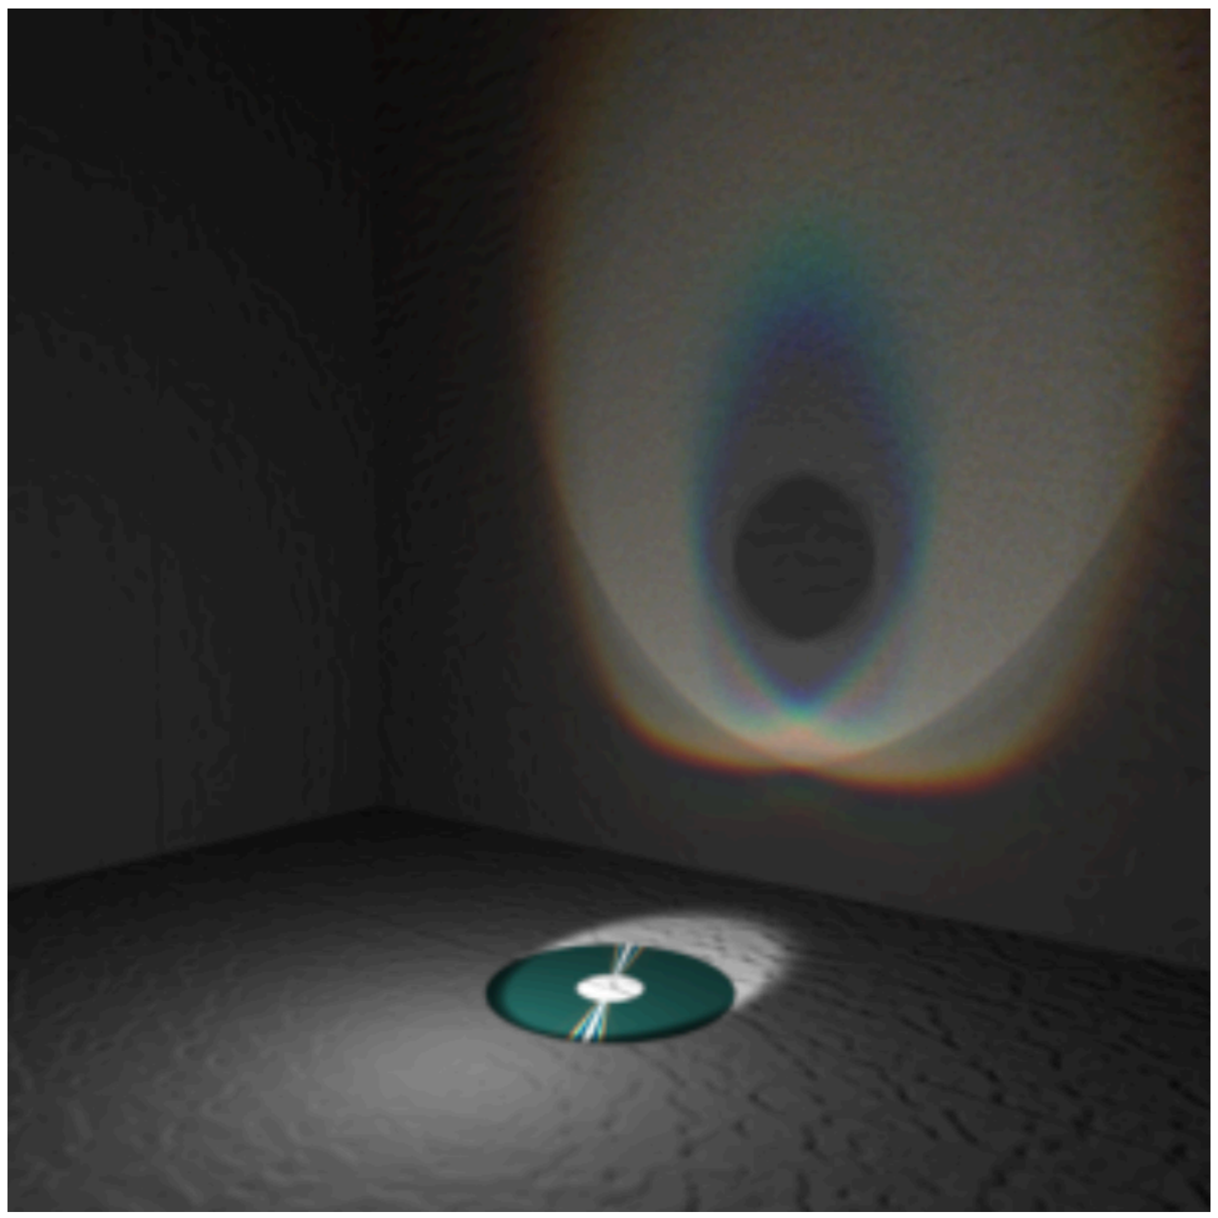

compact disk (CD)
[Cuypers 11]

metallic film
[Laurent 17]

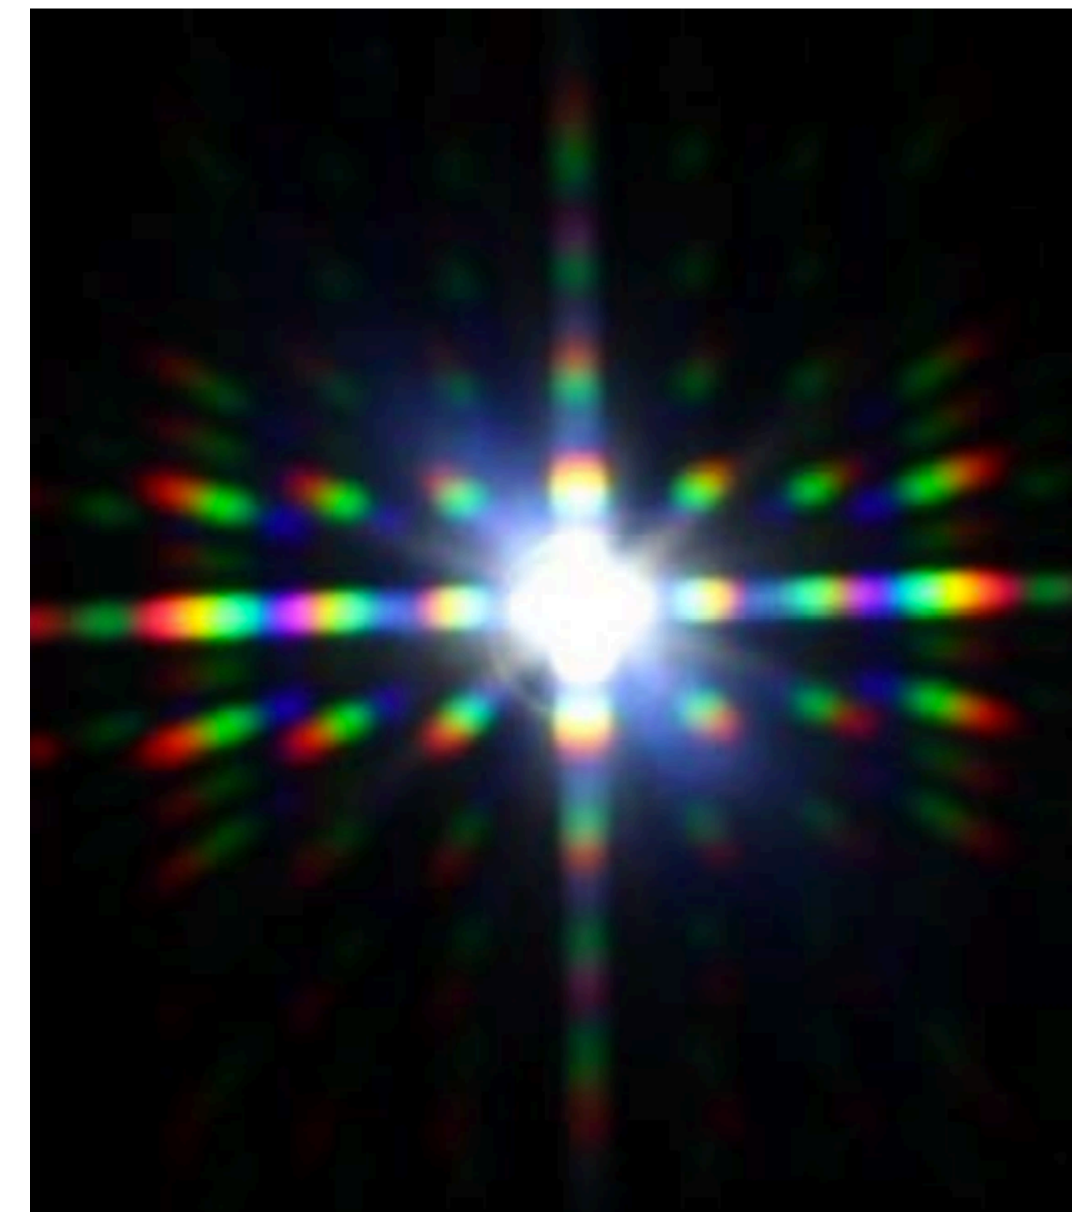

phone screen
[Toisoul 17]

Observations

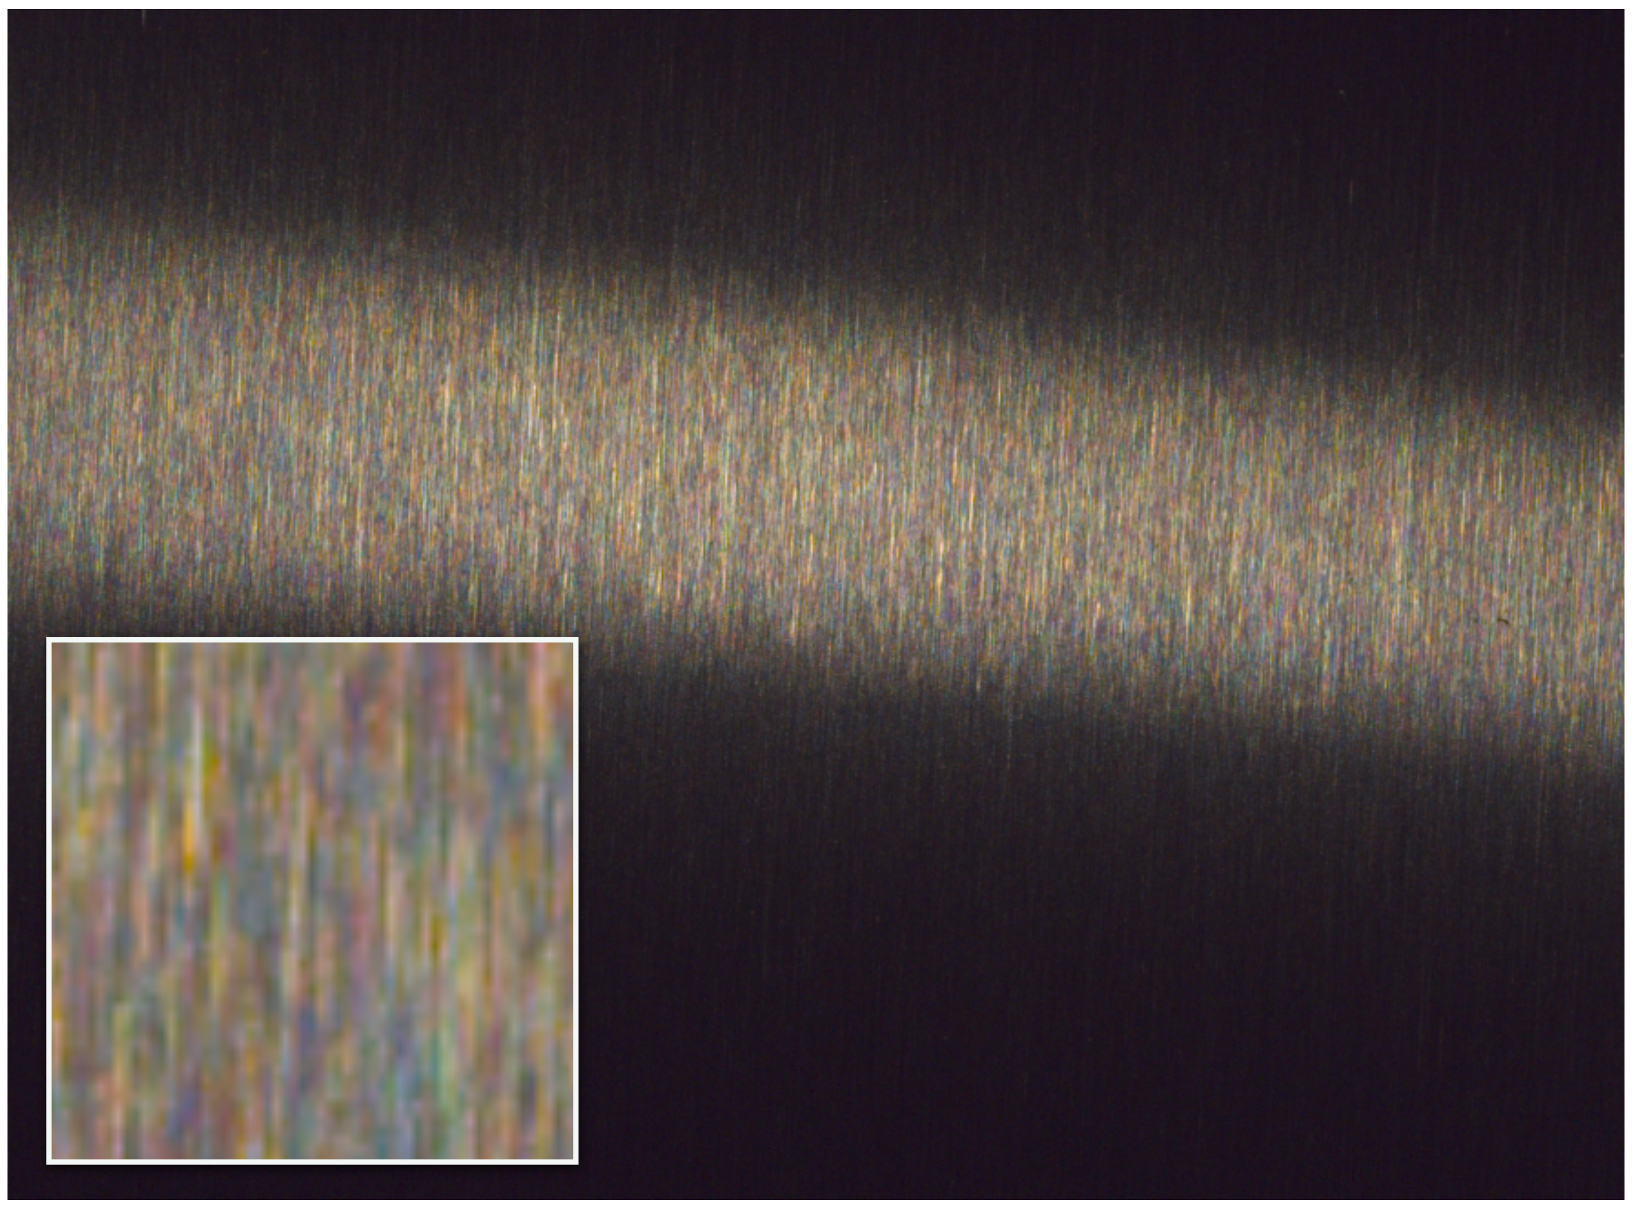

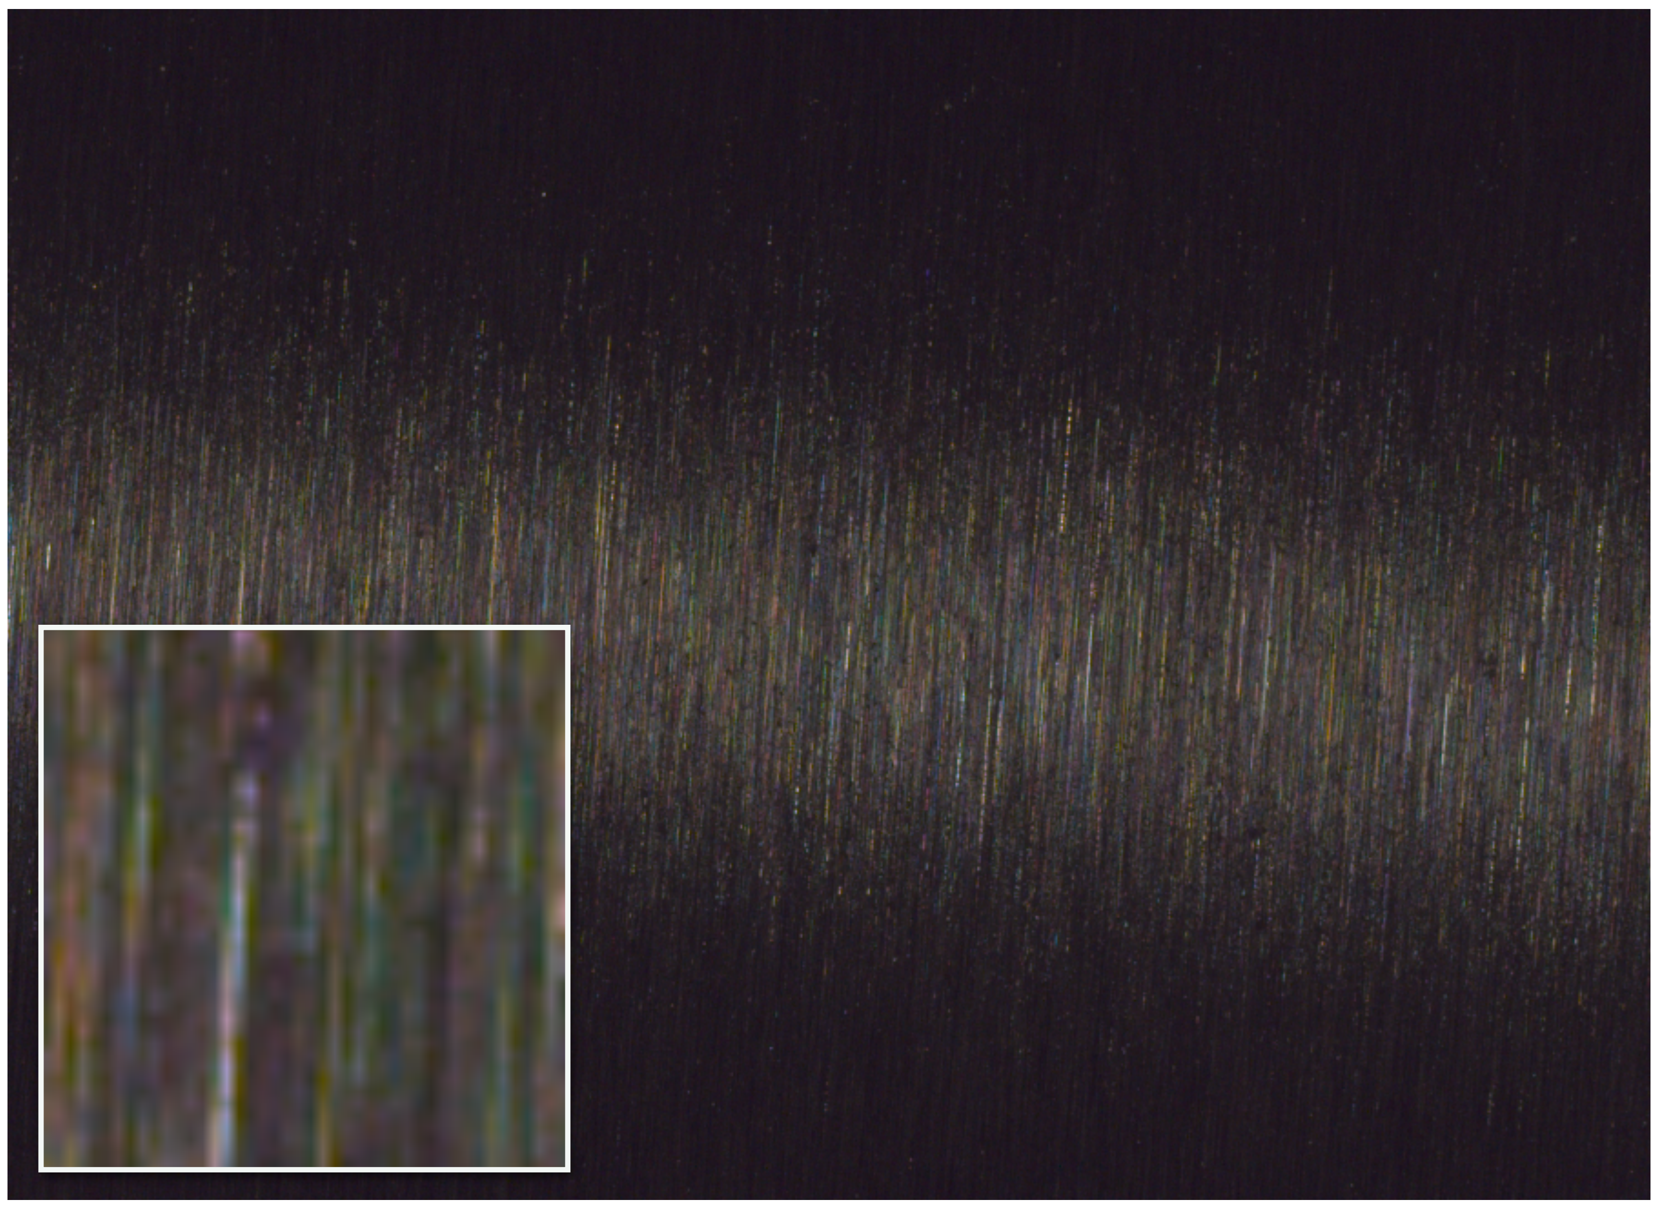

photos of scratched metal

Observations

photo of a Macbook

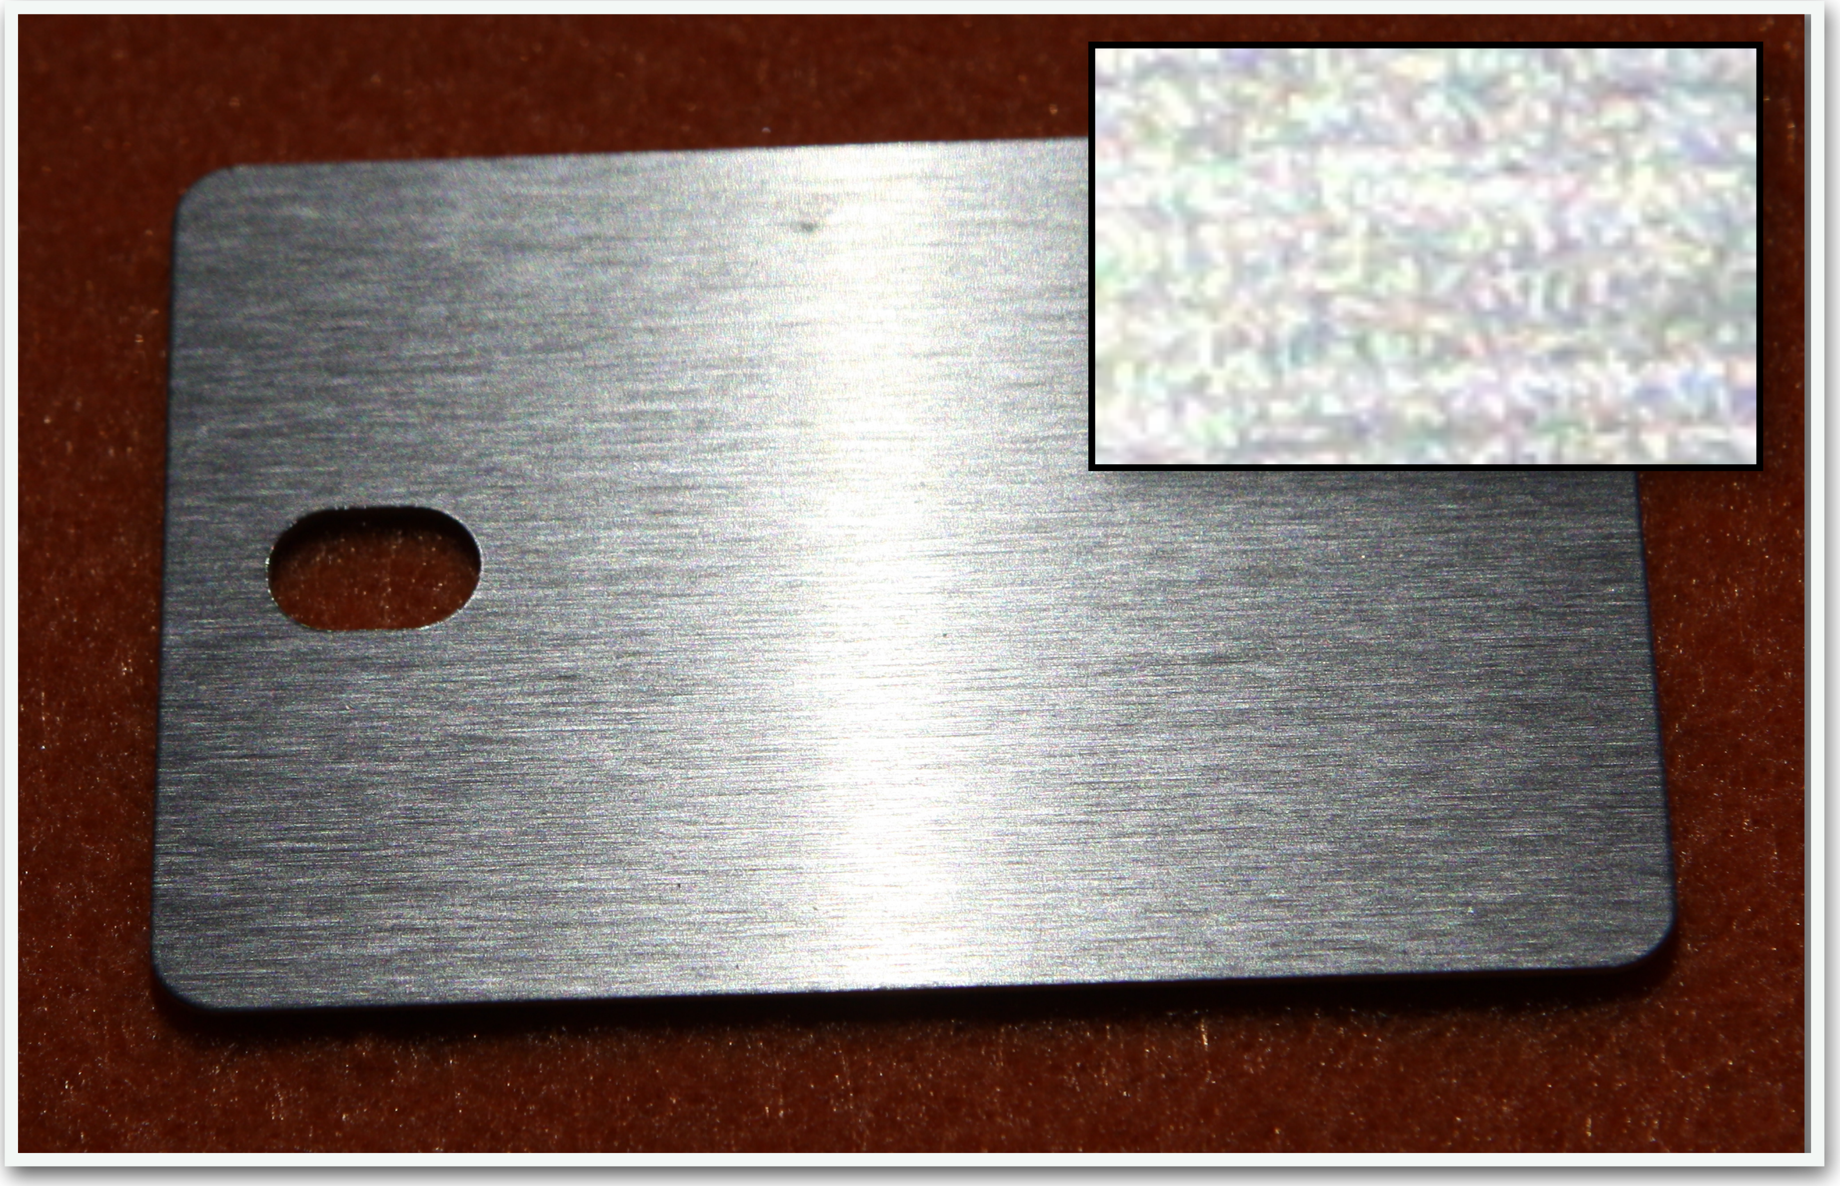

photo of an aluminum patch

Detailed Material under Wave Optics

Heightfields

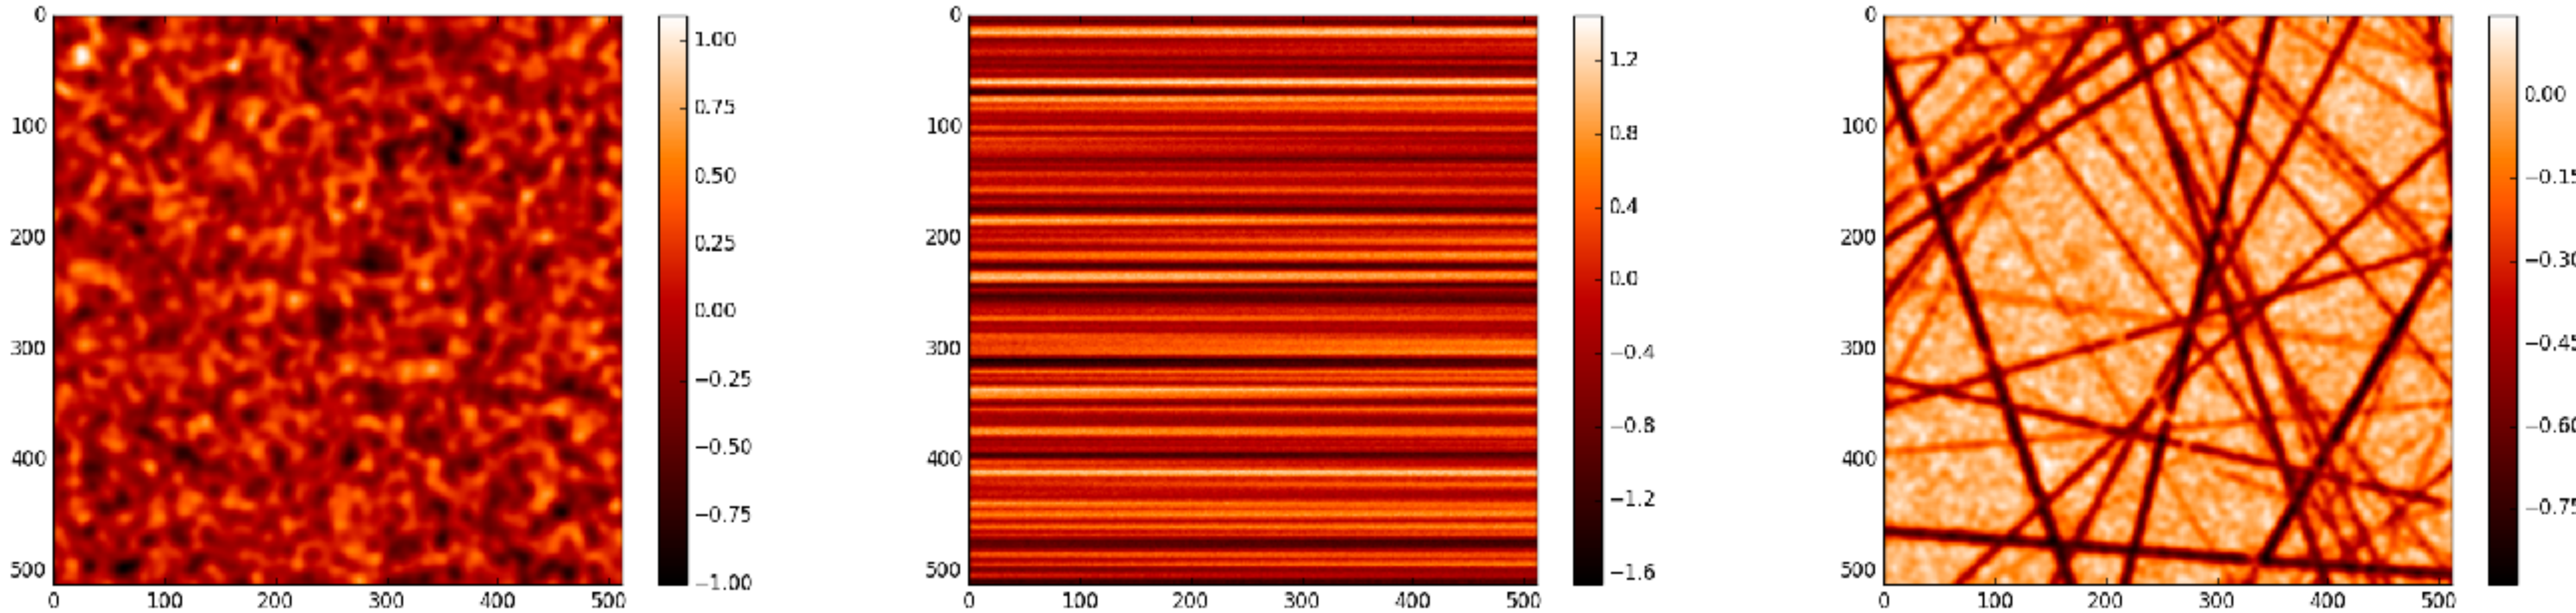

isotropic **brushed** **scratched**

BRDFs

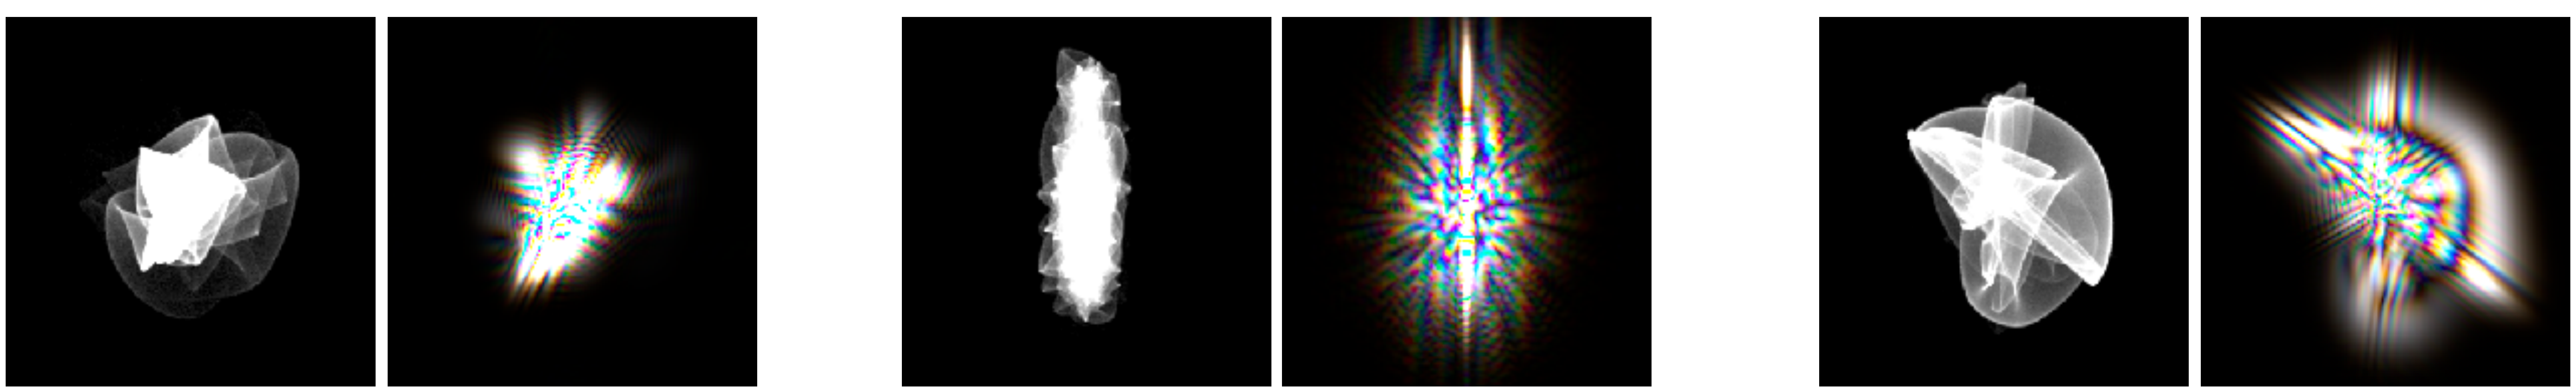

geometric **wave** **geometric** **wave** **geometric** **wave**

Detailed Material under Wave Optics

[Yan et al. 2018]


fn
control
alt
option
⌘
command


fn
control
alt
option
⌘
command

What is it about?

Procedural Appearance

Procedural Appearance

- Can we define details without textures?
  - Yes! Compute a noise function on the fly.

[Lagae et al. 2009]

Procedural Appearance

- Can we define details without textures?
  - Yes! Compute a noise function on the fly.
  - 3D noise -> internal structure if cut or broken

[Lagae et al. 2009]

Procedural Appearance

- Can we define details without textures?
  - Yes! Compute a noise function on the fly.
  - Thresholding
    (noise -> binary noise)

Example:

```
if noise(x, y, z) > threshold:
    reflectance = 1
else:
    reflectance = 0
```

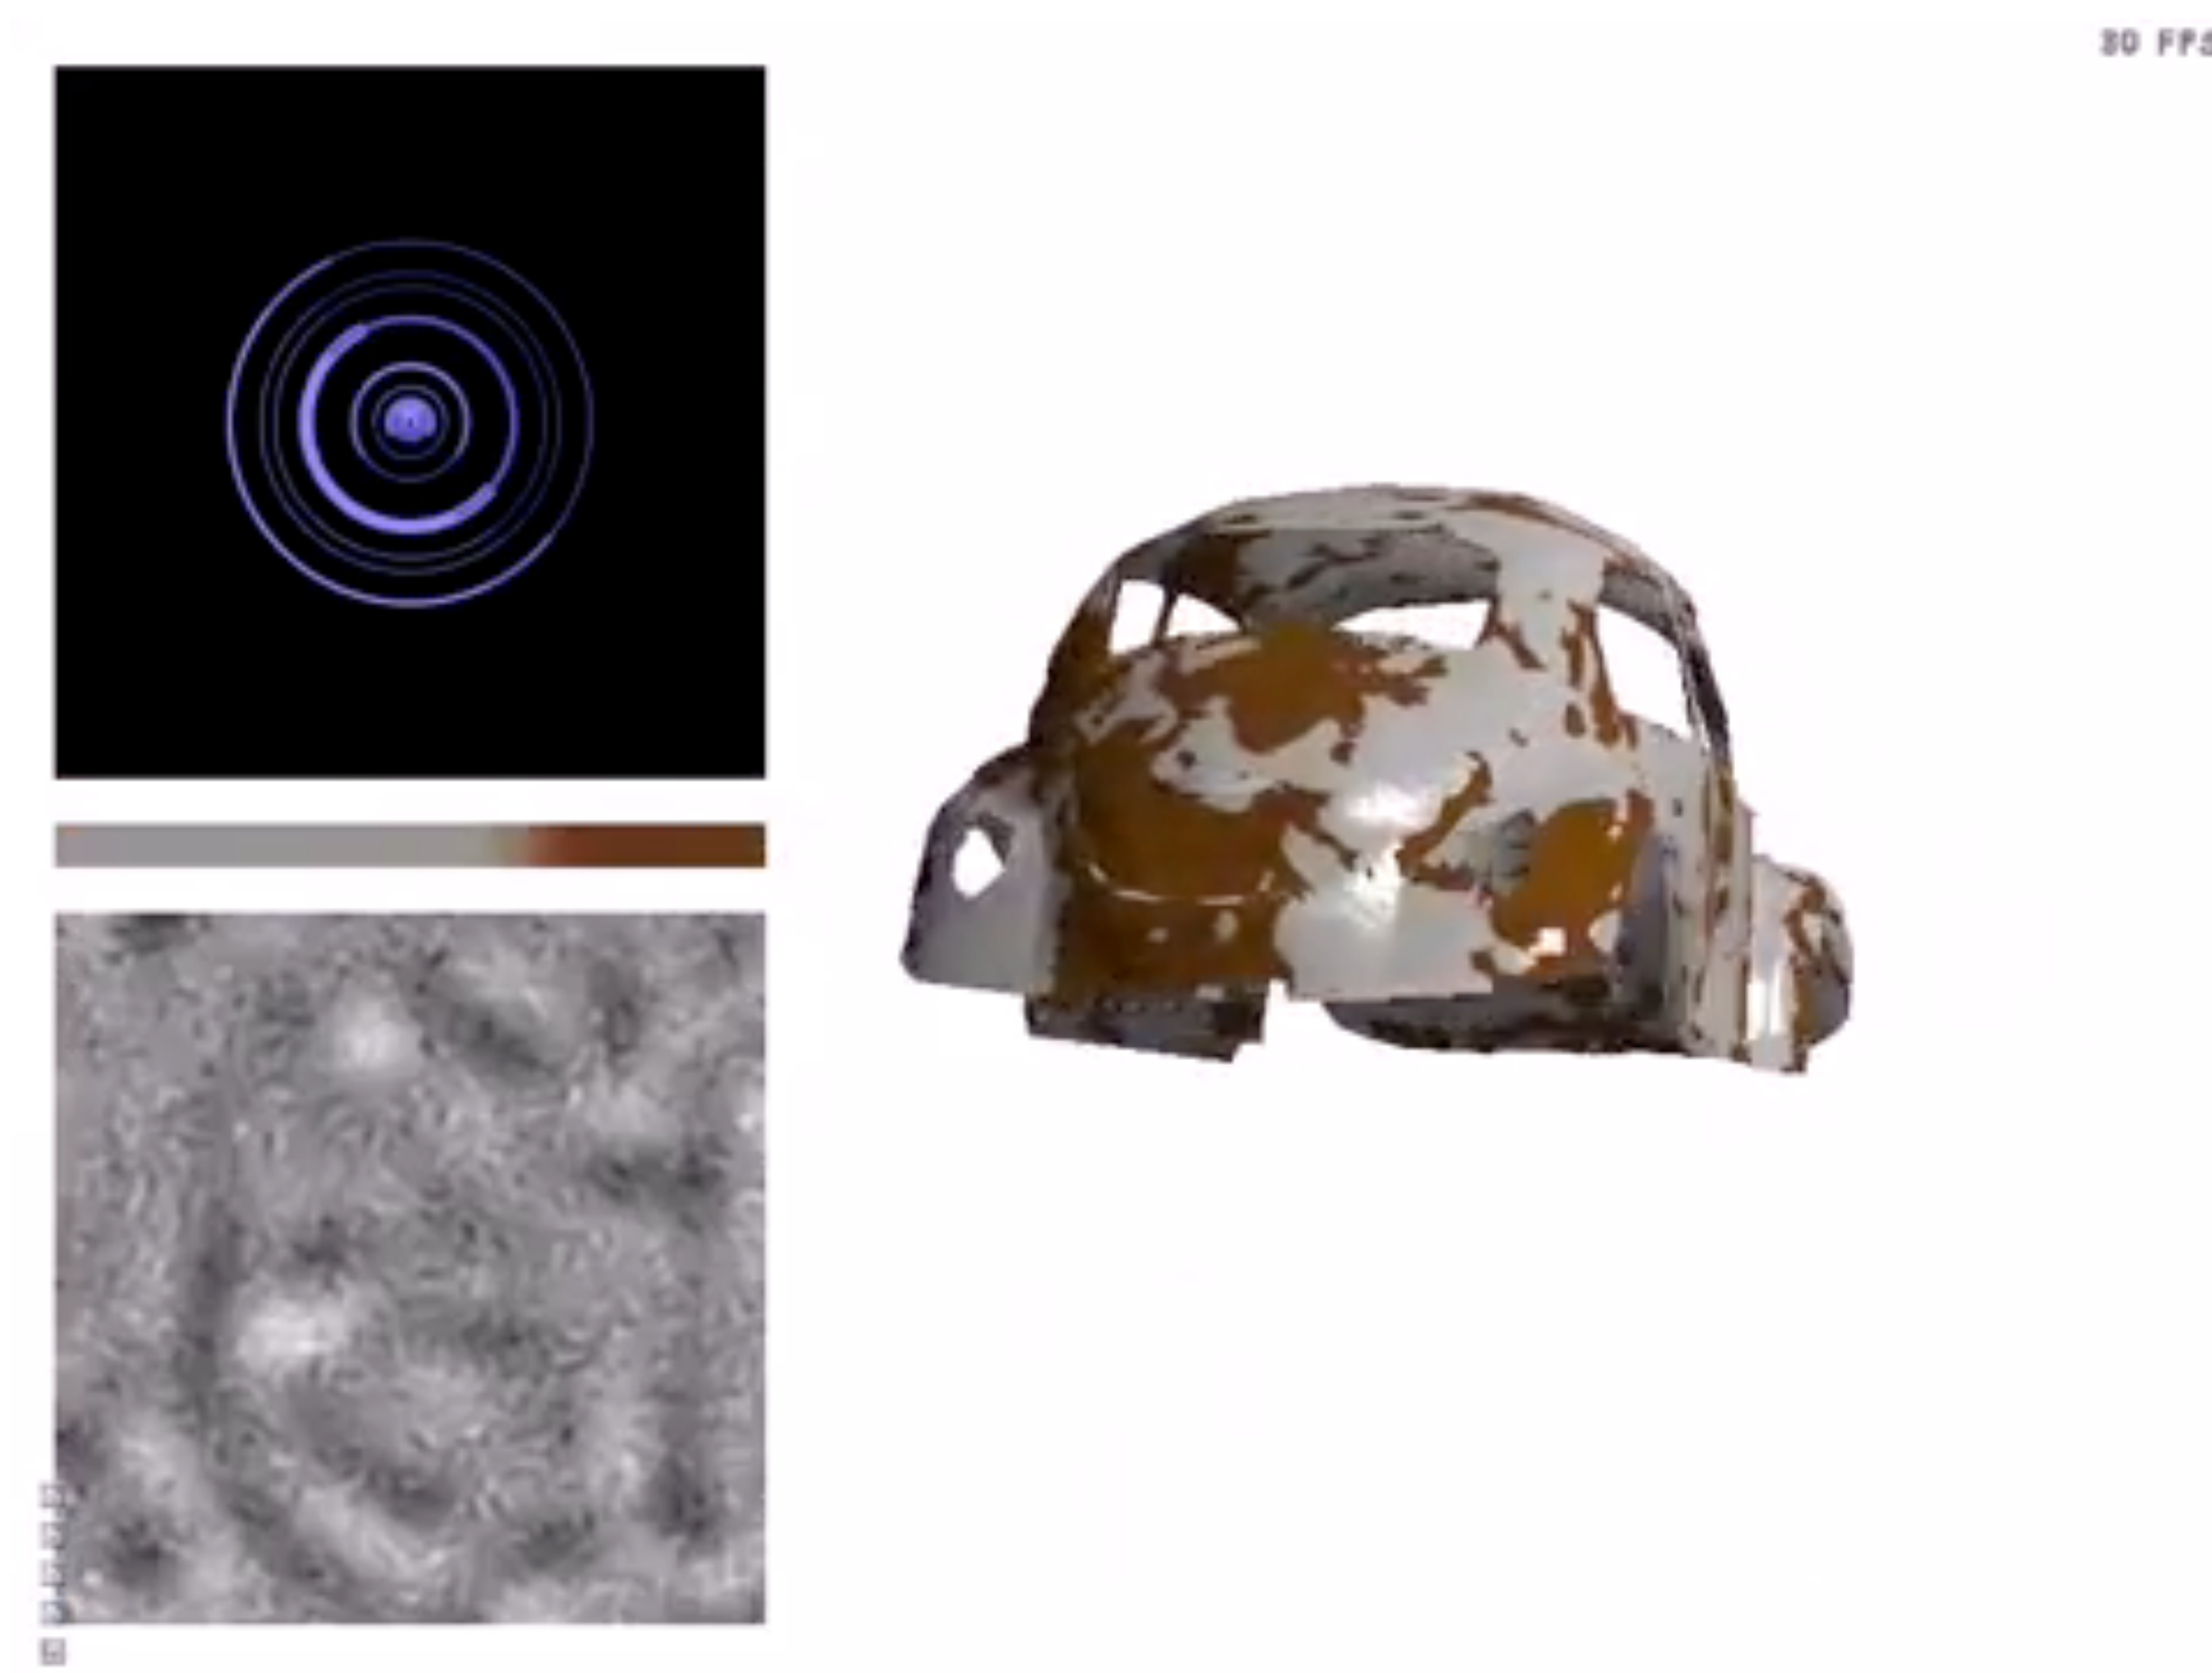

[Lagae et al. 2009]

Procedural Appearance

- Complex noise functions can be very powerful.

Procedural Appearance

- Complex noise functions can be very powerful.

[Steve Worley]

Procedural Appearance

- Complex noise functions can be very powerful.

[Liu et al. 2016]

Procedural Appearance

- Complex noise functions can be very powerful.

[Liu et al. 2016]

Thank you!

(And thank Prof. Ren Ng for many of the slides!)