Lecture 22: Light and shading

Announcements

- PS10 out
- 2nd-to-last lecture on low-level vision.
- Rest of course: recent vision topics.

Many interpretations of color!

The Workshop Metaphor

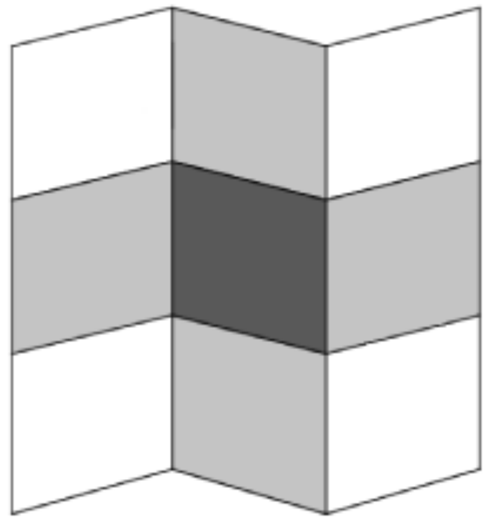

(a) an image

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Source: J. Barron

The Workshop Metaphor

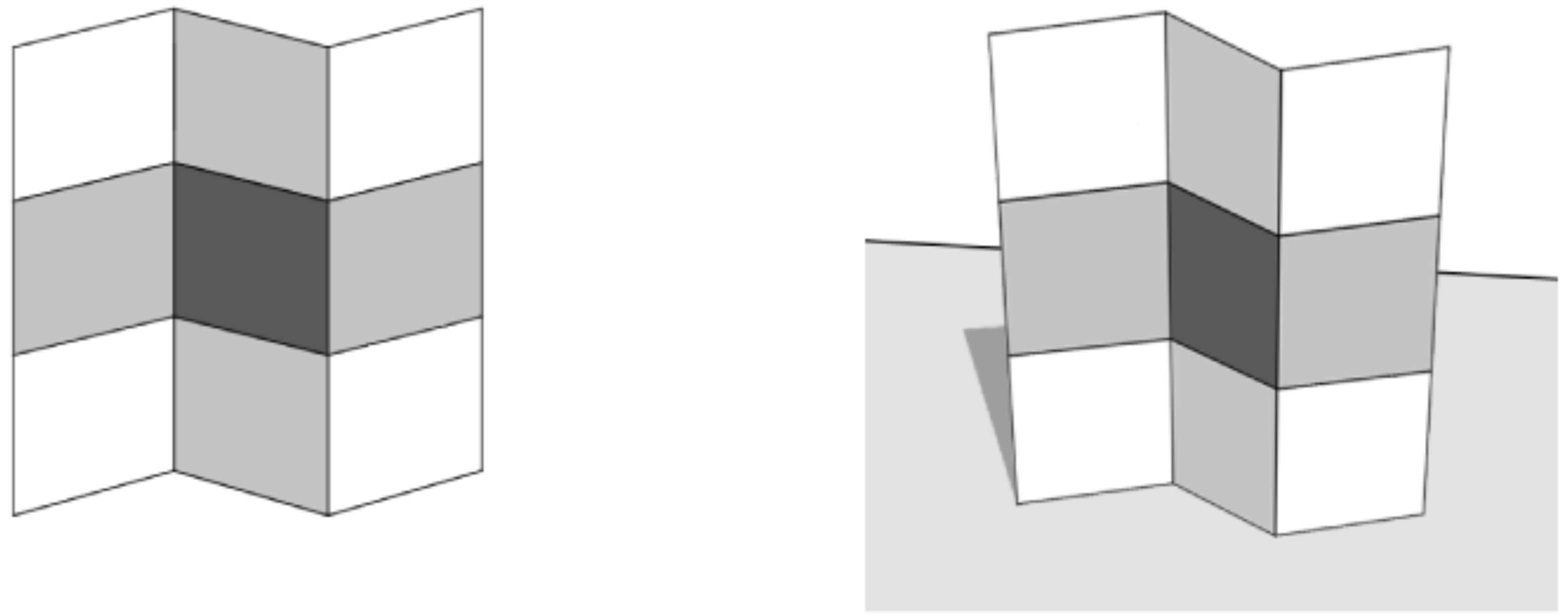

(a) an image (b) a likely explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Source: J. Barron

The Workshop Metaphor

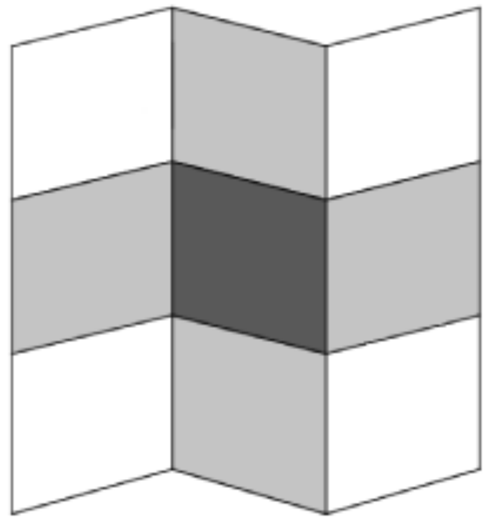

(a) an image

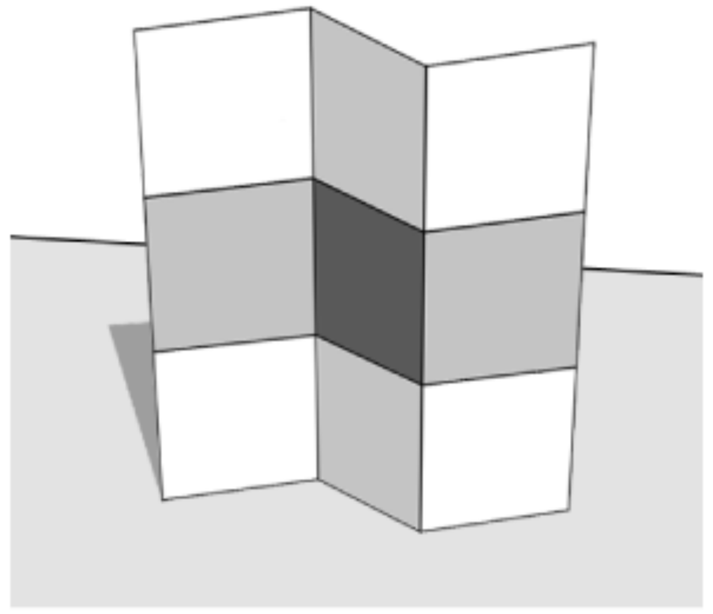

(b) a likely explanation

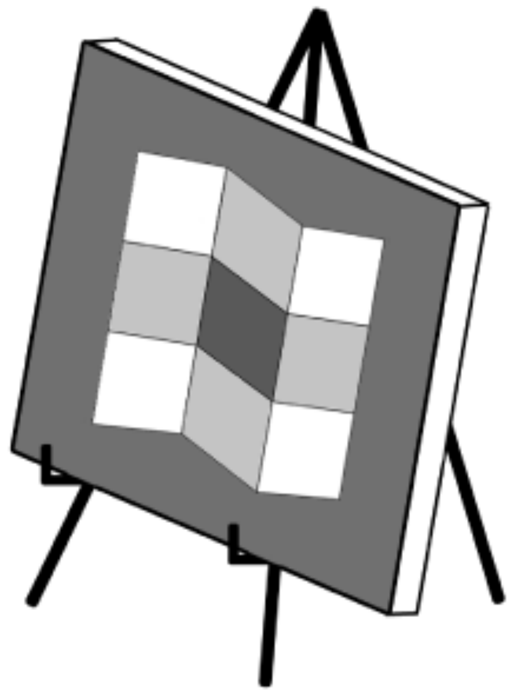

(c) painter's explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Source: J. Barron

The Workshop Metaphor

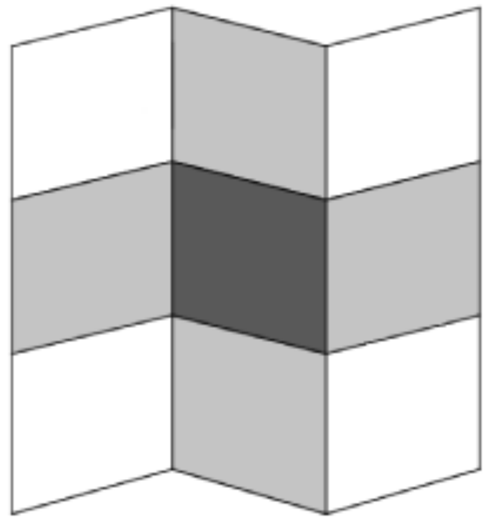

(a) an image

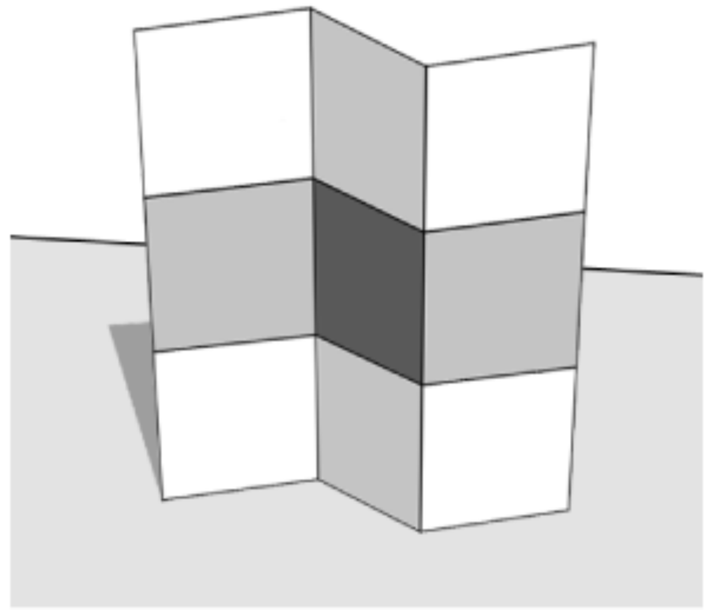

(b) a likely explanation

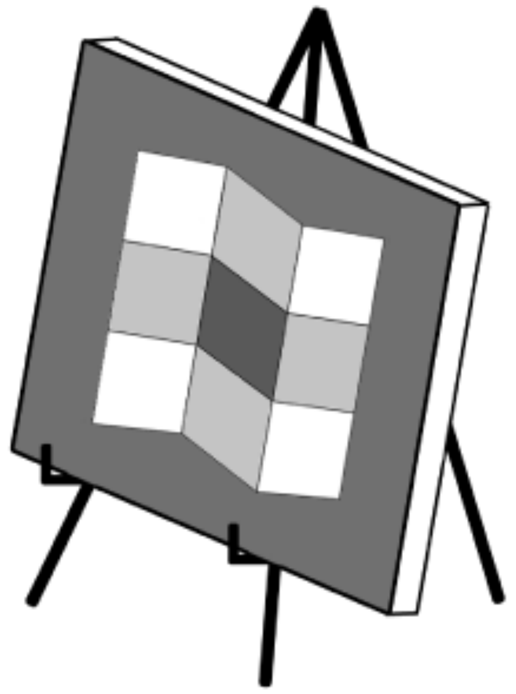

(c) painter's explanation

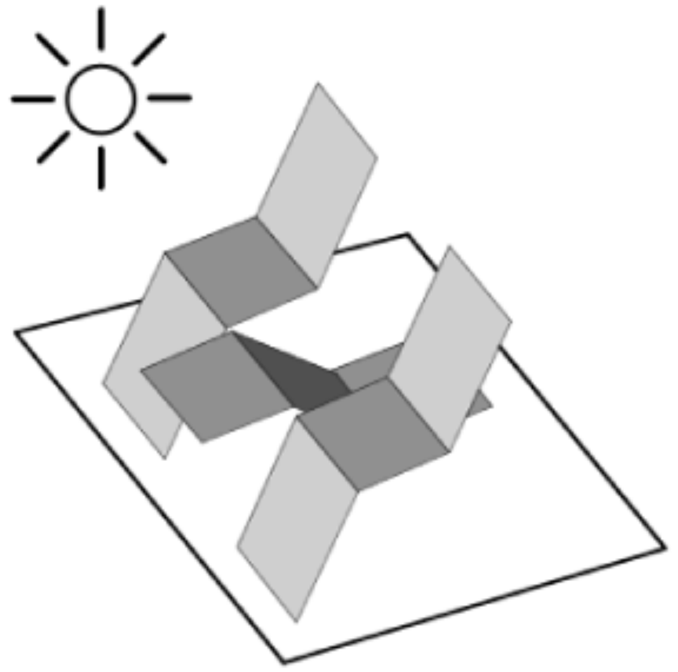

(d) sculptor's explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Source: J. Barron

The Workshop Metaphor

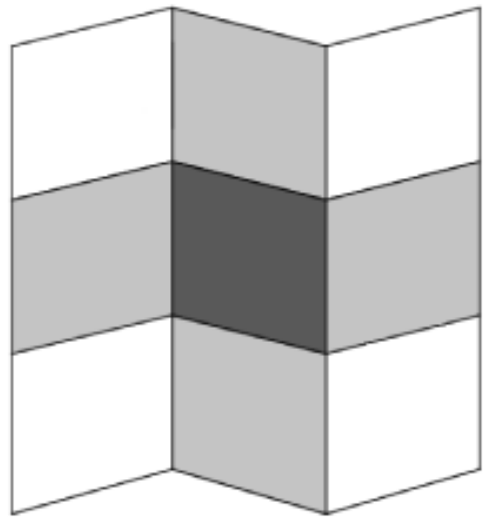

(a) an image

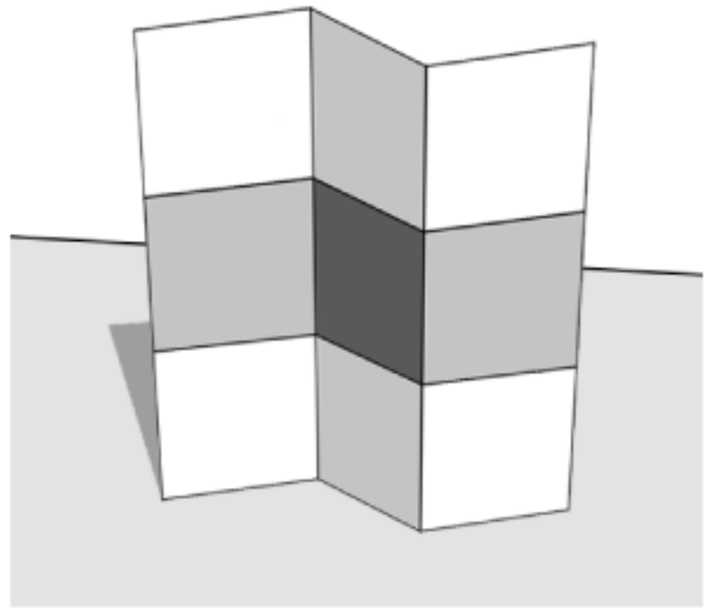

(b) a likely explanation

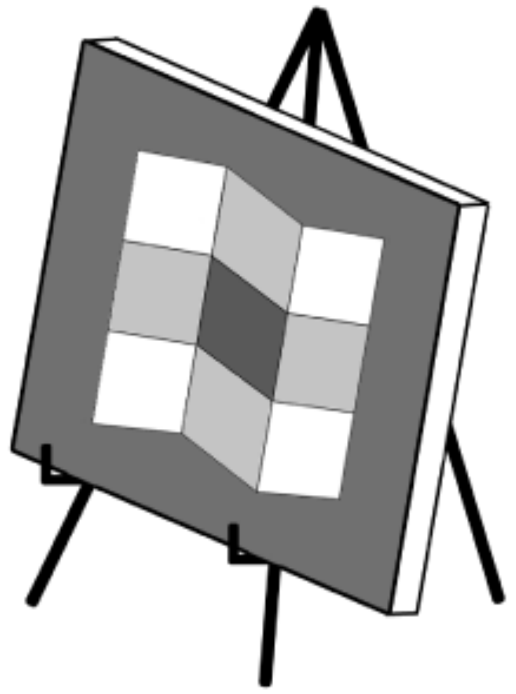

(c) painter's explanation

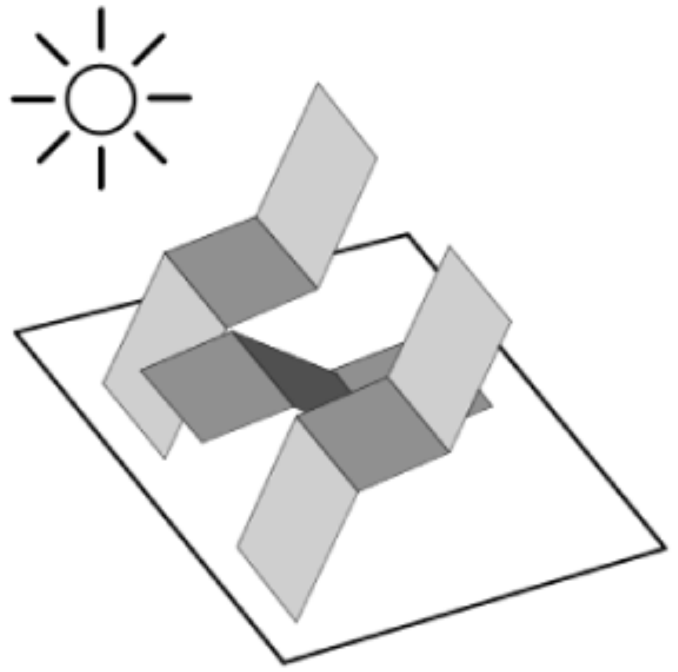

(d) sculptor's explanation

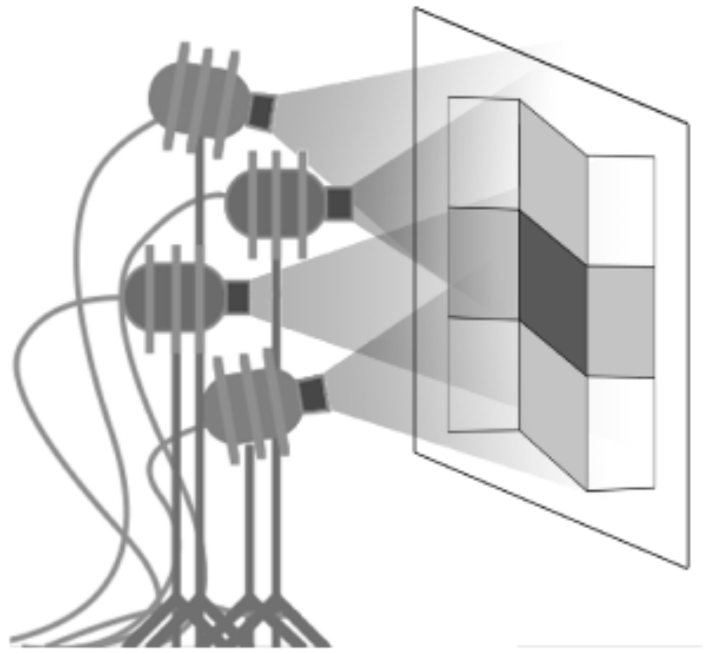

(e) gaffer's explanation

E. Adelson and A. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996.

Source: J. Barron

Today

- Light and surfaces
- Shape from shading
- Photometric stereo
- Intrinsic image decomposition

Recall: interaction of light and surfaces

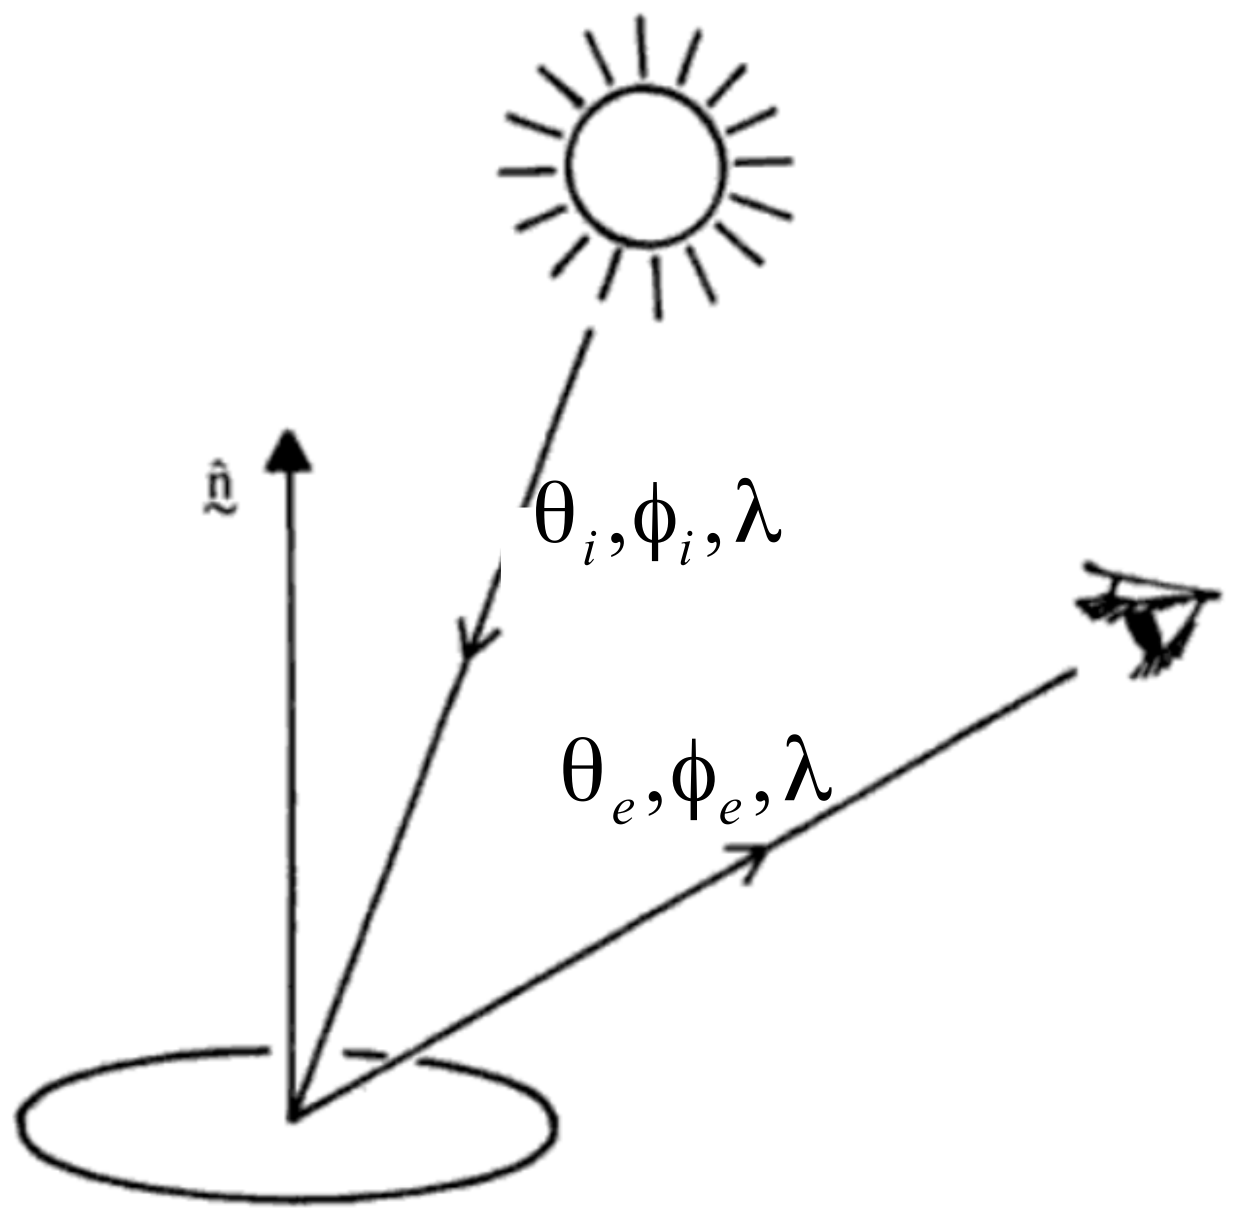

Figure 10-7. The bidirectional reflectance distribution function is the ratio of the radiance of the surface patch as viewed from the direction (θ_e, ϕ_e) to the irradiance resulting from illumination from the direction (θ_i, ϕ_i).

[Horn, 1986]

Spectral radiance: power in a specified direction, per unit area, per unit solid angle, per unit wavelength.

$$BRDF = f(\theta_i, \phi_i, \theta_e, \phi_e, \lambda) = \frac{L(\theta_e, \phi_e, \lambda)}{E(\theta_i, \phi_i, \lambda)}$$

Spectral irradiance: incident power per unit area, per unit wavelength

Source: W. Freeman

For now, ignore specular reflection

Source: Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra. Changes by N. Snavely

And Refraction…

Source: Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra. Changes by N. Snavely

And Interreflections…

Source: Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra. Changes by N. Snavely

Recall: effect of BRDF on sphere rendering

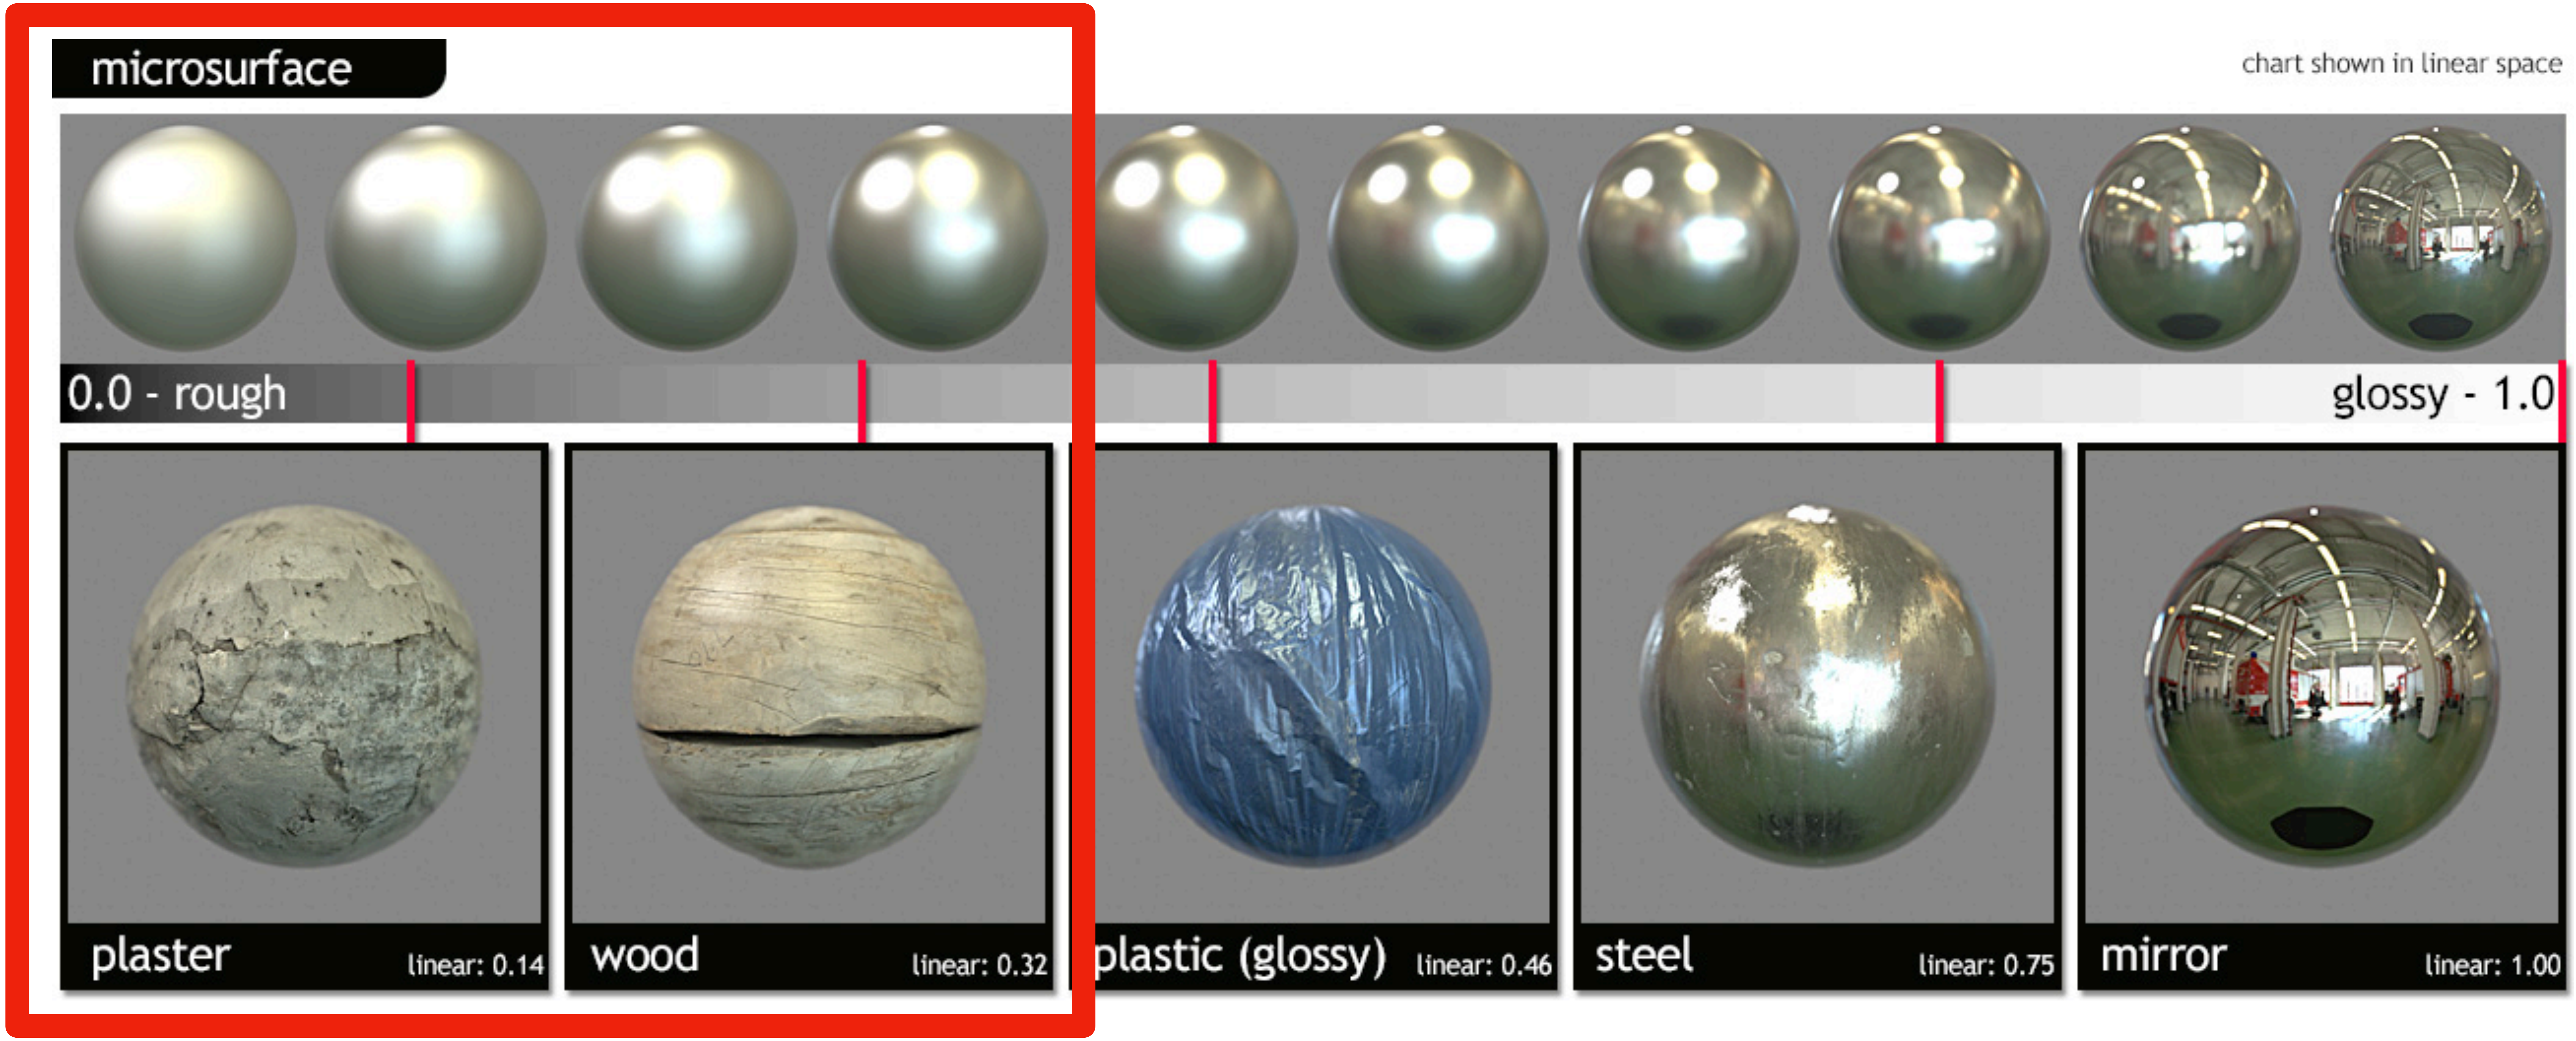

Diffuse/Lambertian reflection

https://marmoset.co/posts/physically-based-rendering-and-you-can-too/

Source: W. Freeman

Diffuse reflection

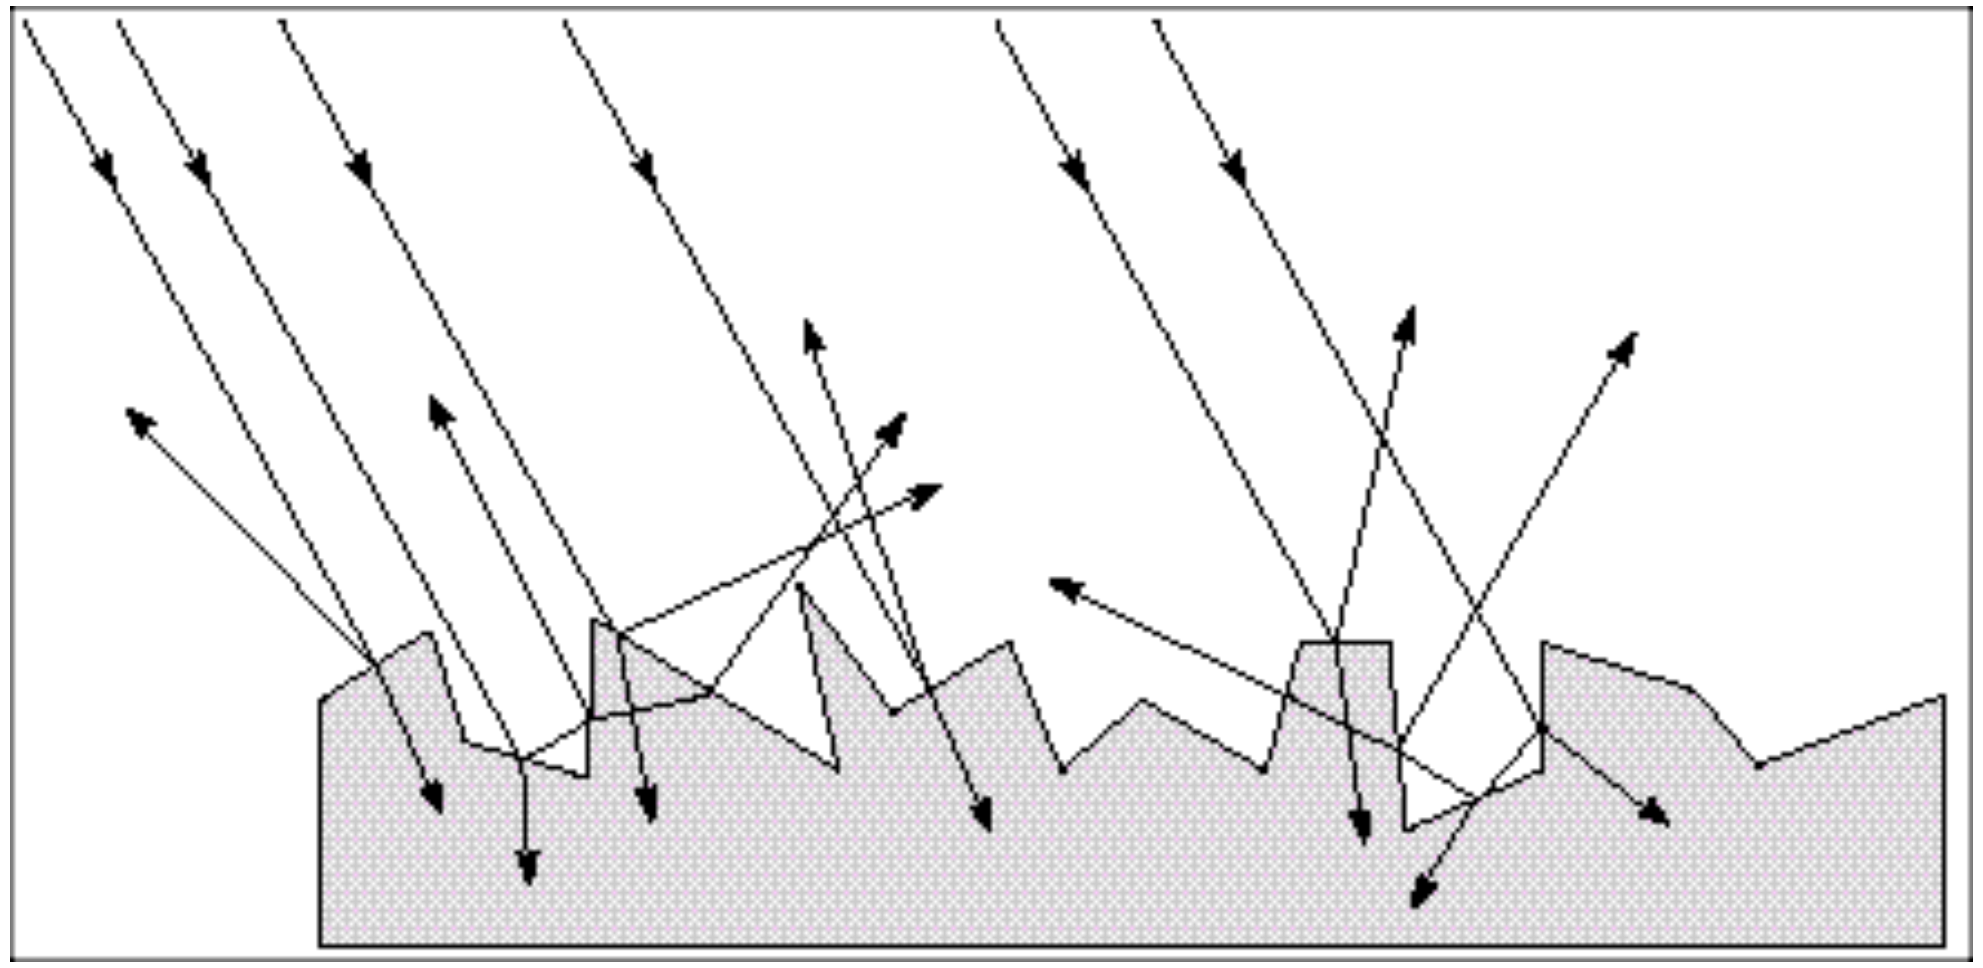

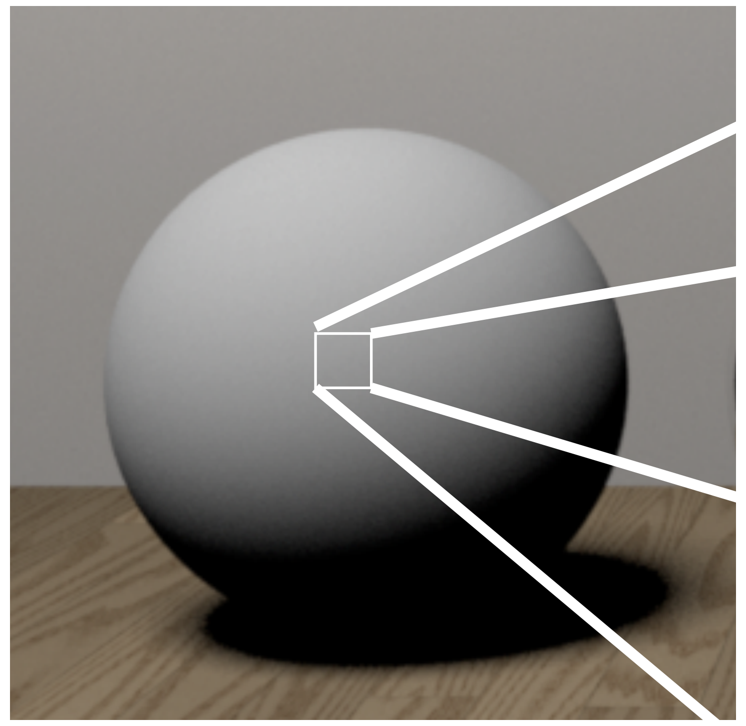

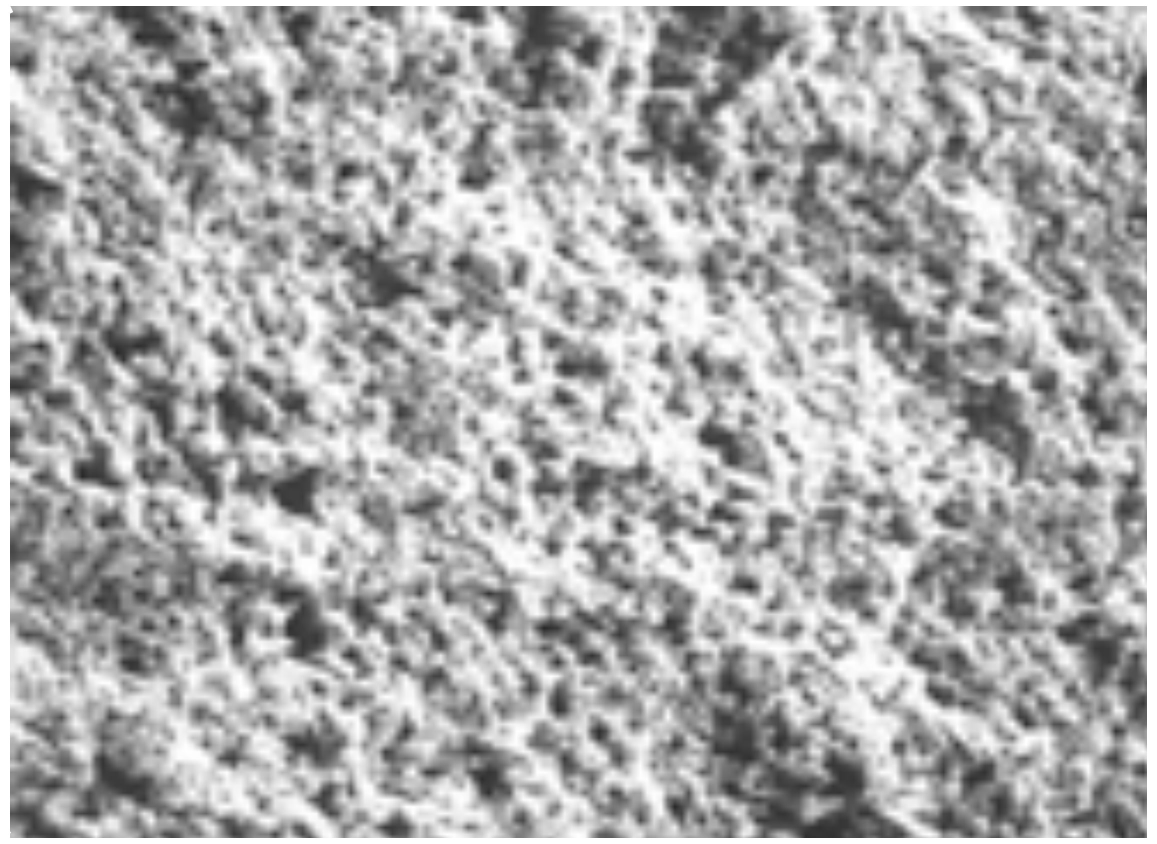

Diffuse reflection

- Dull, matte surfaces like chalk or latex paint
- Microfacets scatter incoming light randomly
- Effect is that light is reflected equally in all directions

Source: S. Lazebnik and K. Bala

Directional lighting

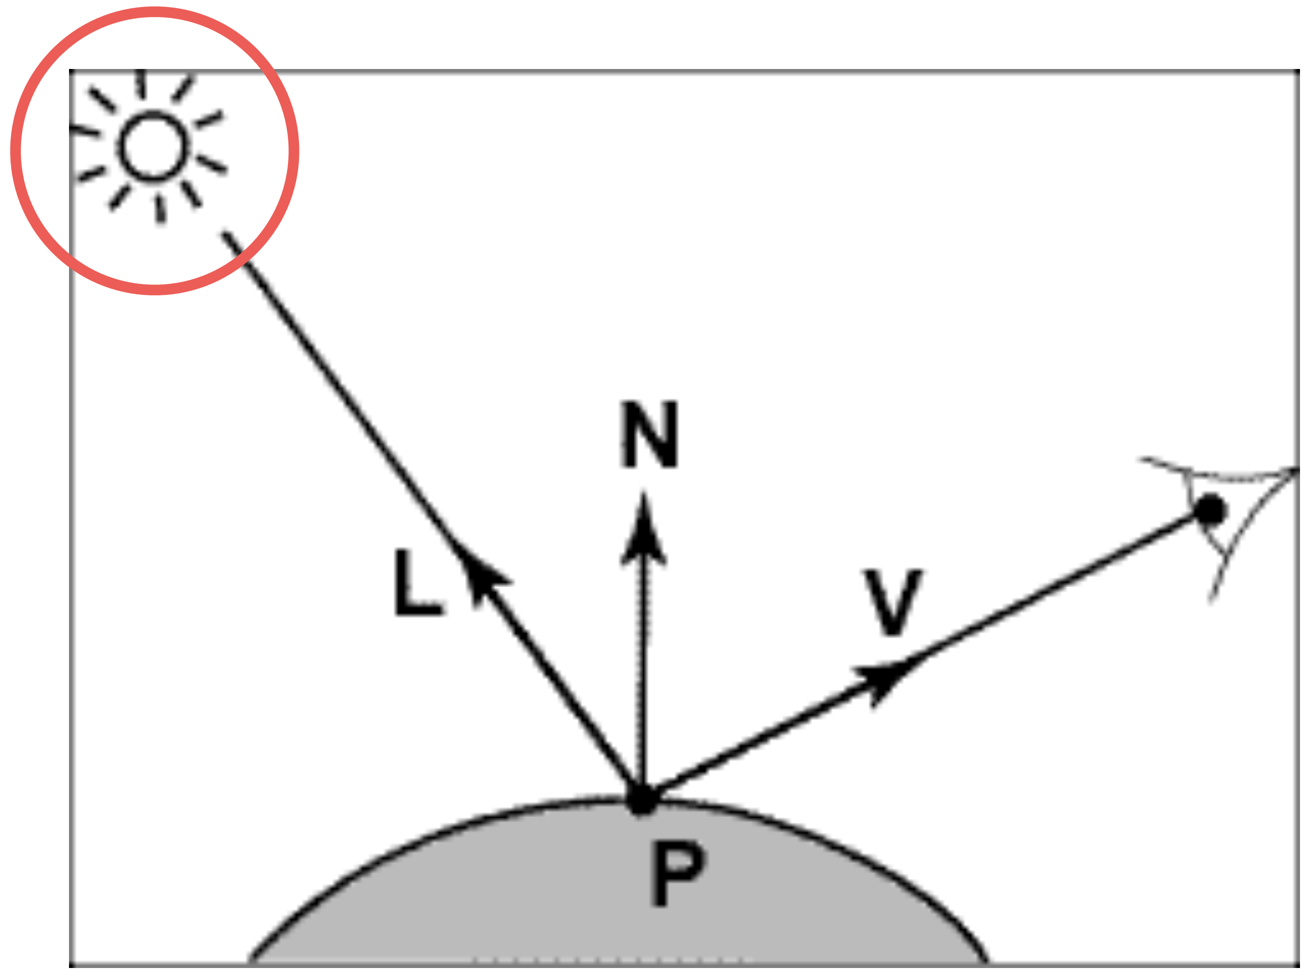

- All rays are parallel
- Equivalent to an infinitely distant point source

Source: N. Snavely

Diffuse reflection

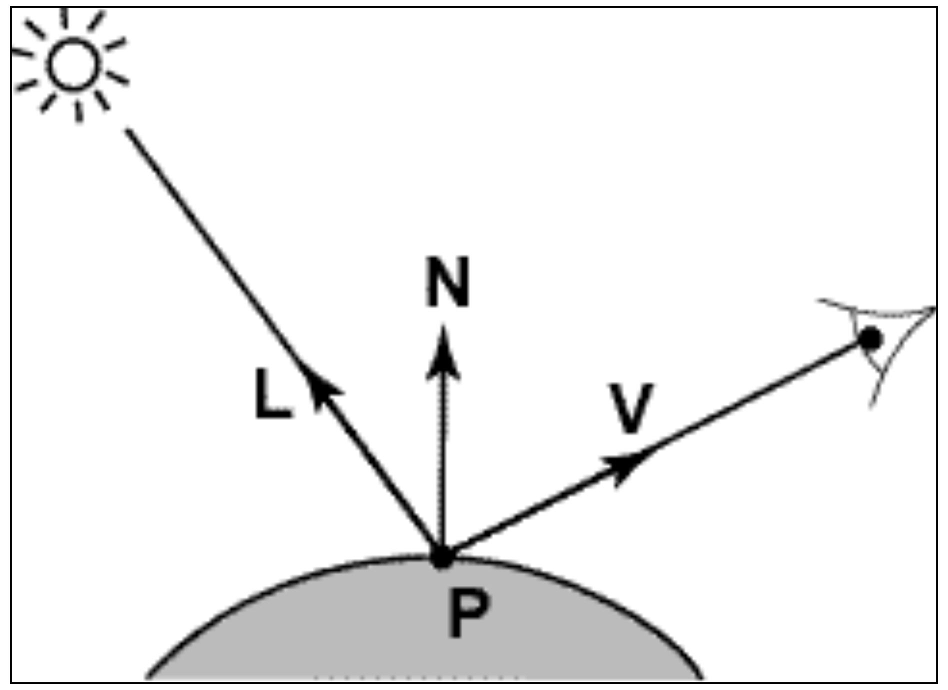

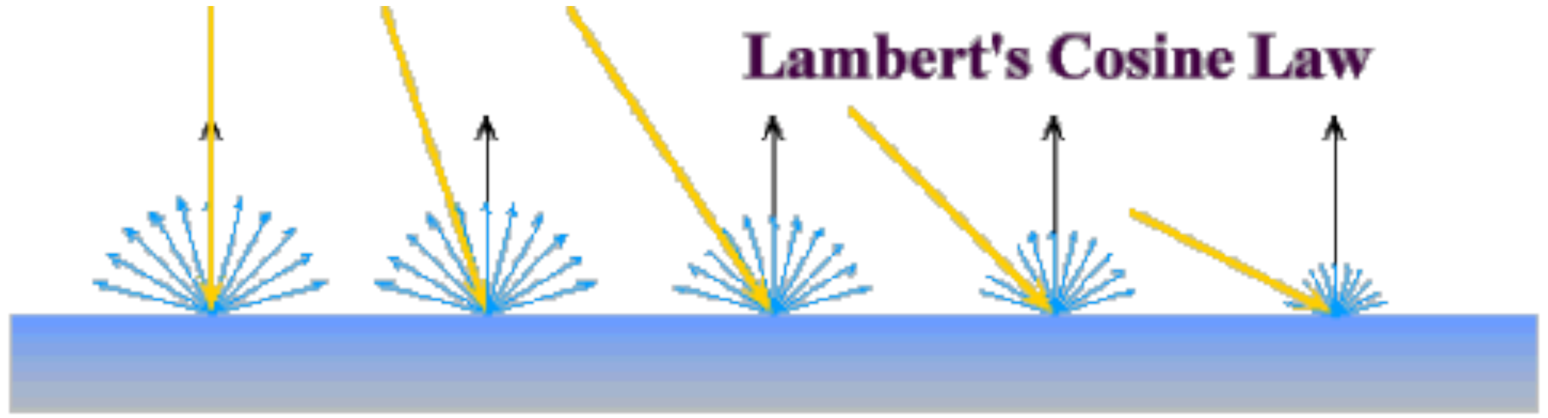

$$R_e = k_d \mathbf{N} \cdot \mathbf{L} R_i$$

image intensity of **P** → $$I = k_d \mathbf{N} \cdot \mathbf{L}$$

Simplifying assumptions we'll often make:

- $I = R_e$: "camera response function" is the identity
  - can always achieve this in practice by inverting it
- $R_i = 1$: light source intensity is 1
  - can achieve this by dividing each pixel in the image by R_i

Source: N. Snavely

Other BRDFs

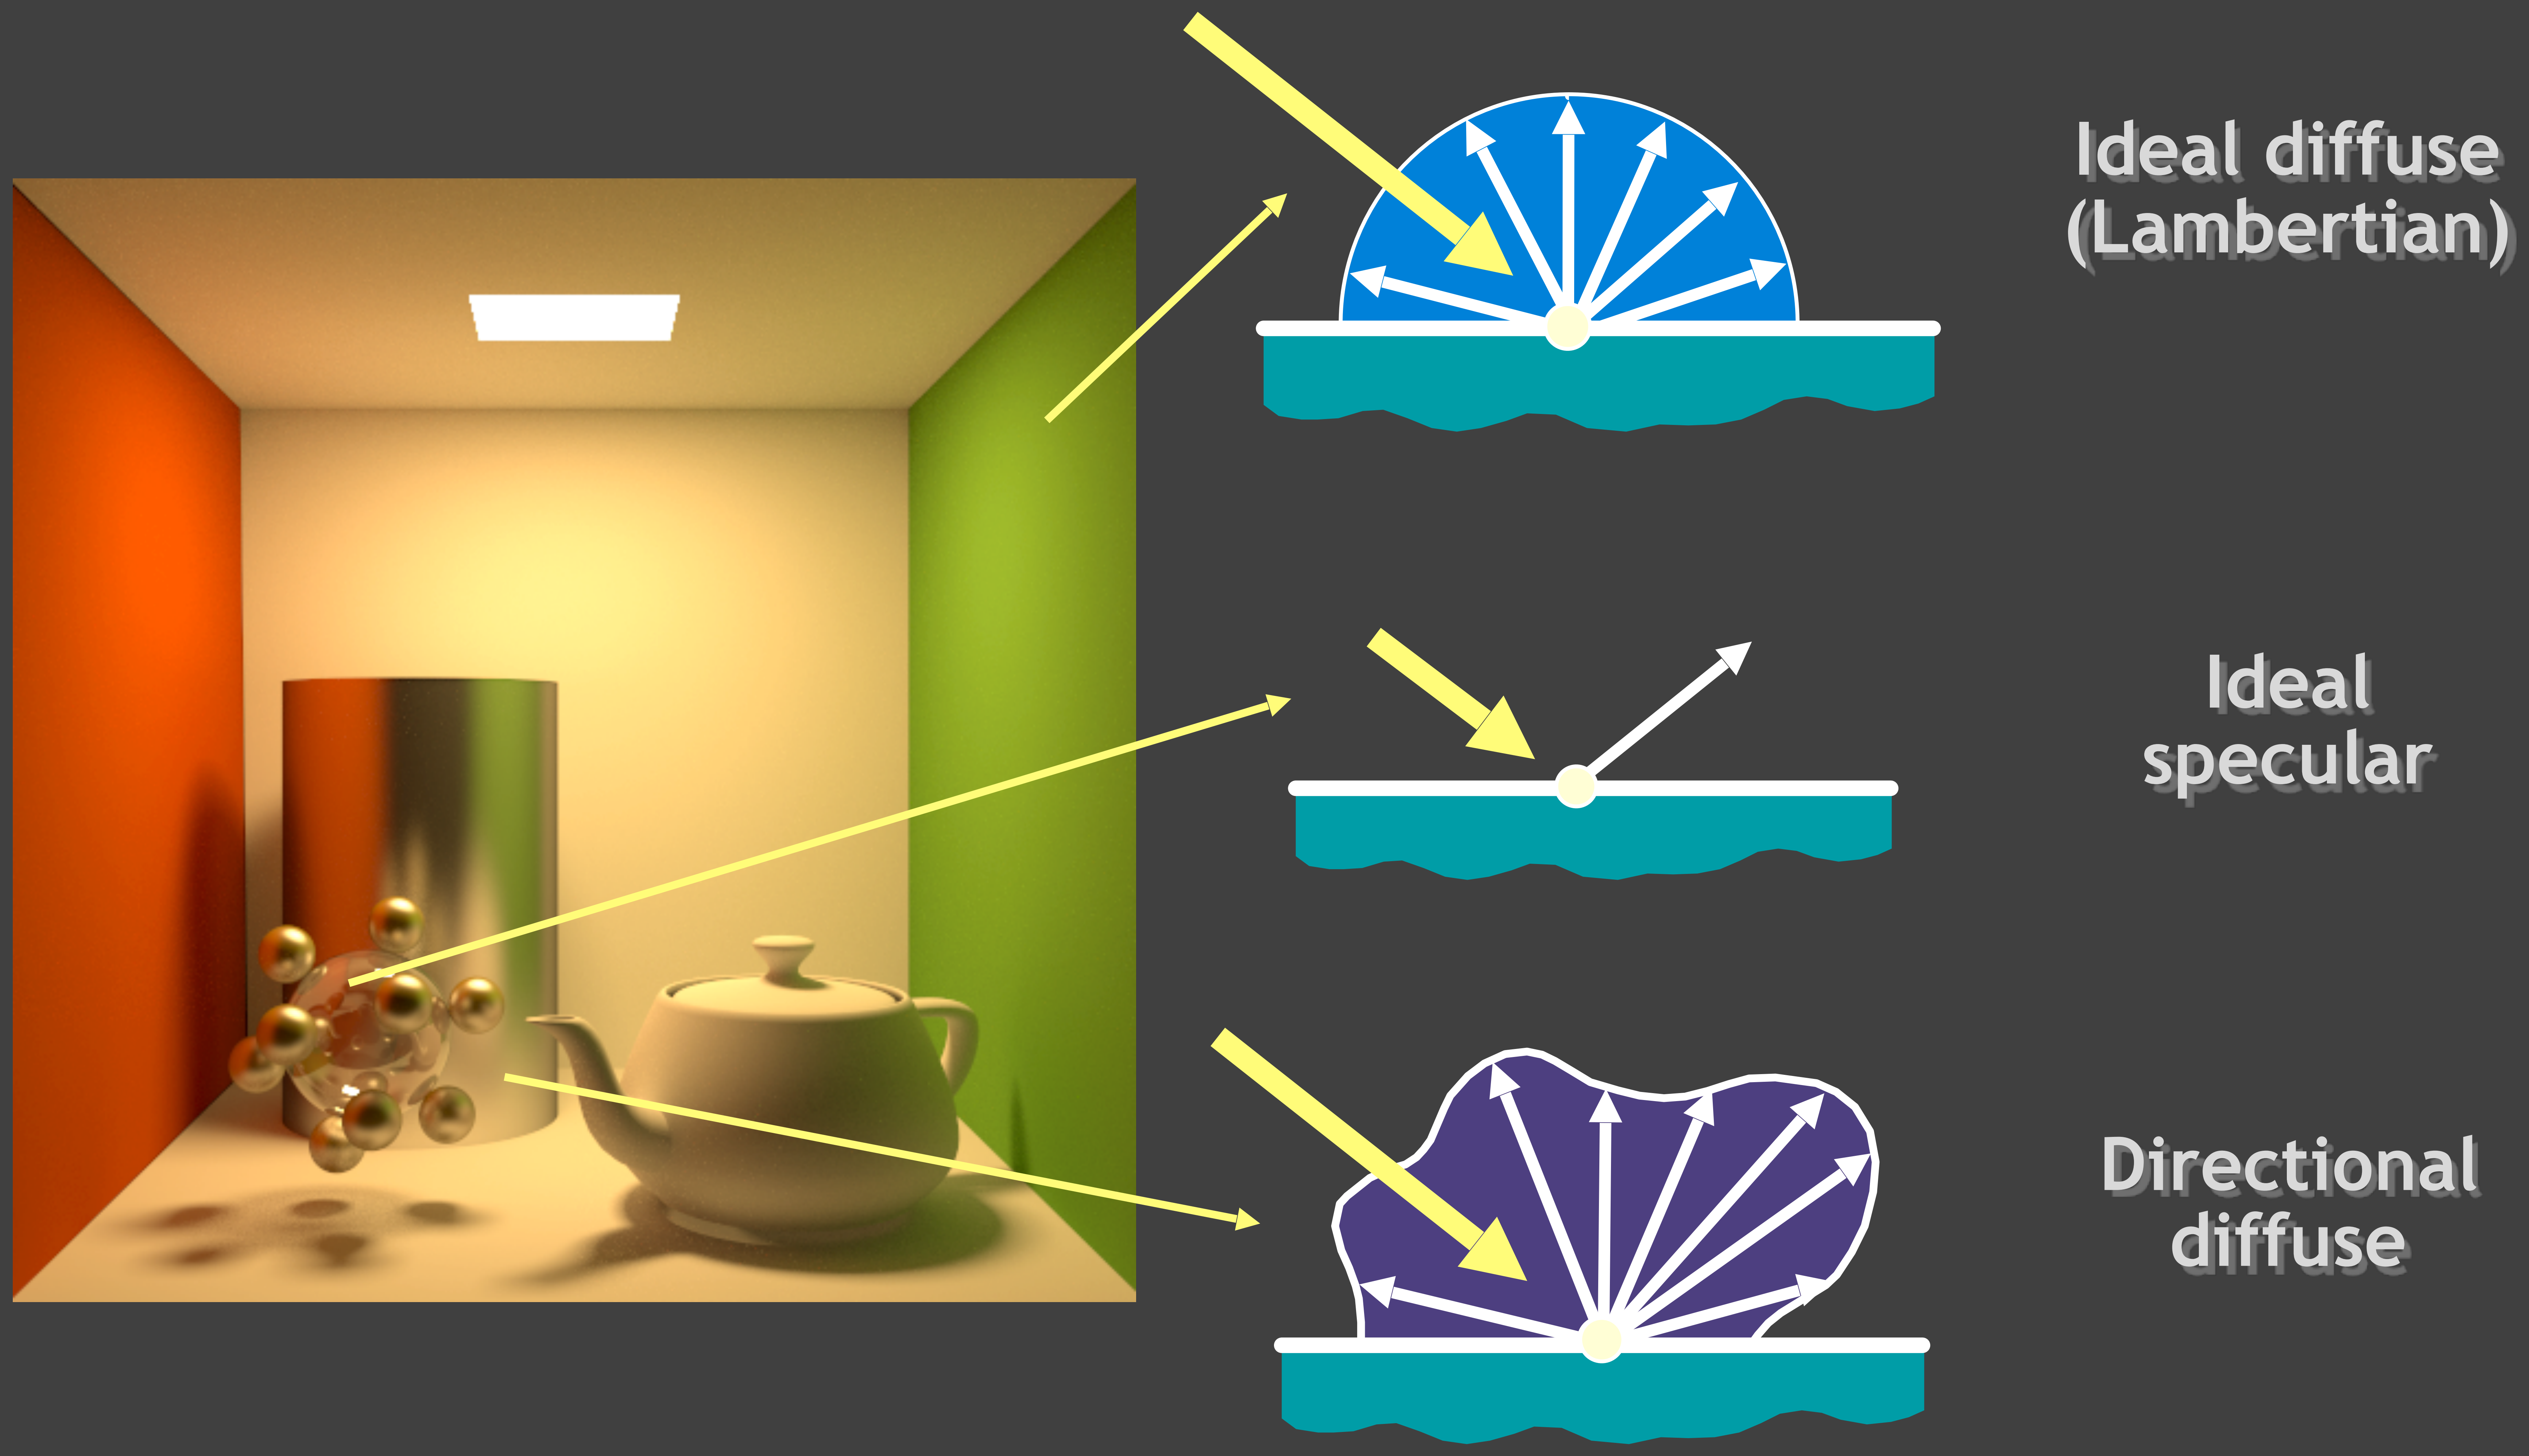
Ideal diffuse
(Lambertian)
Ideal
specular
Directional
diffuse

Non-smooth-surfaced materials

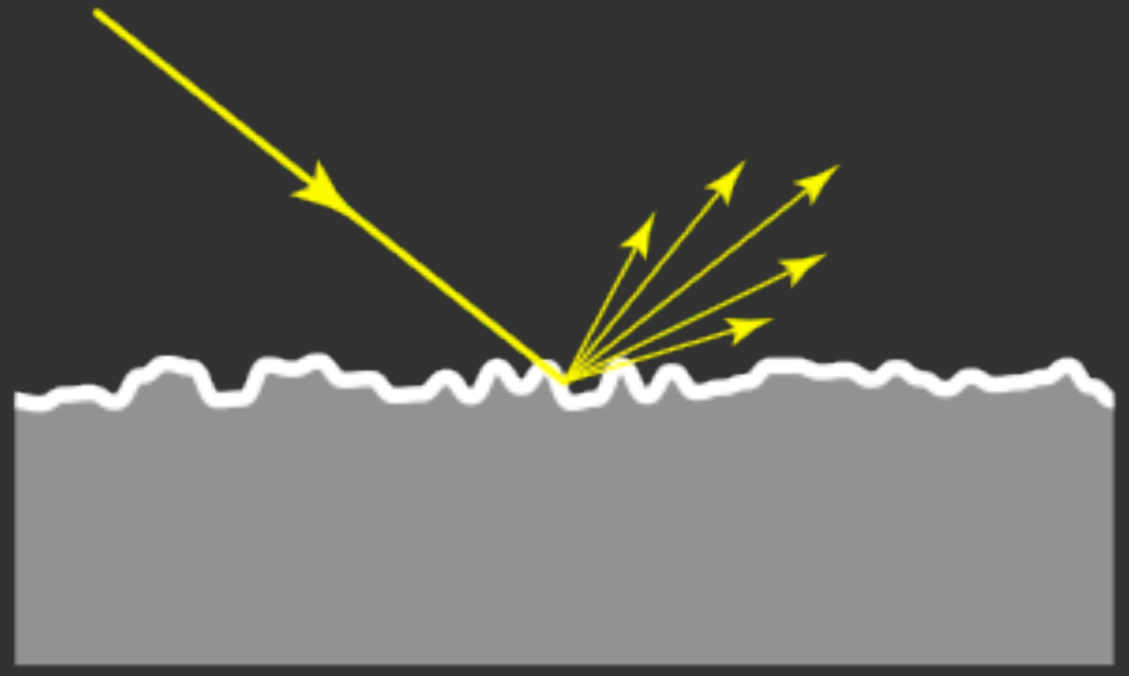

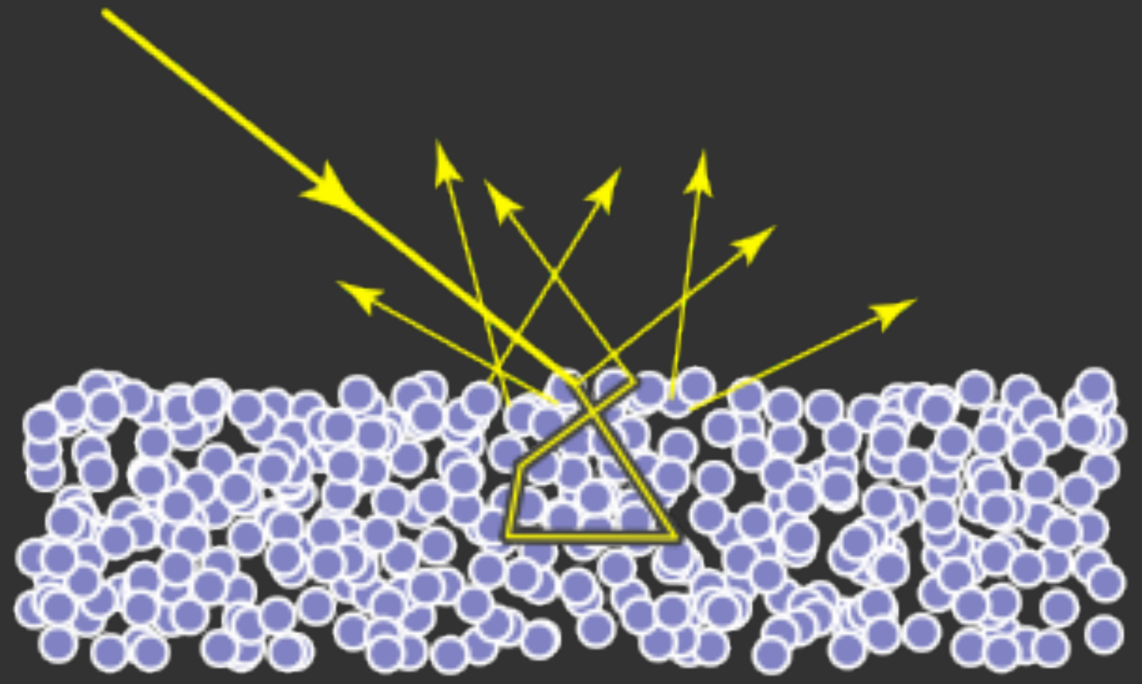

from Steve Marschner

Shape from shading

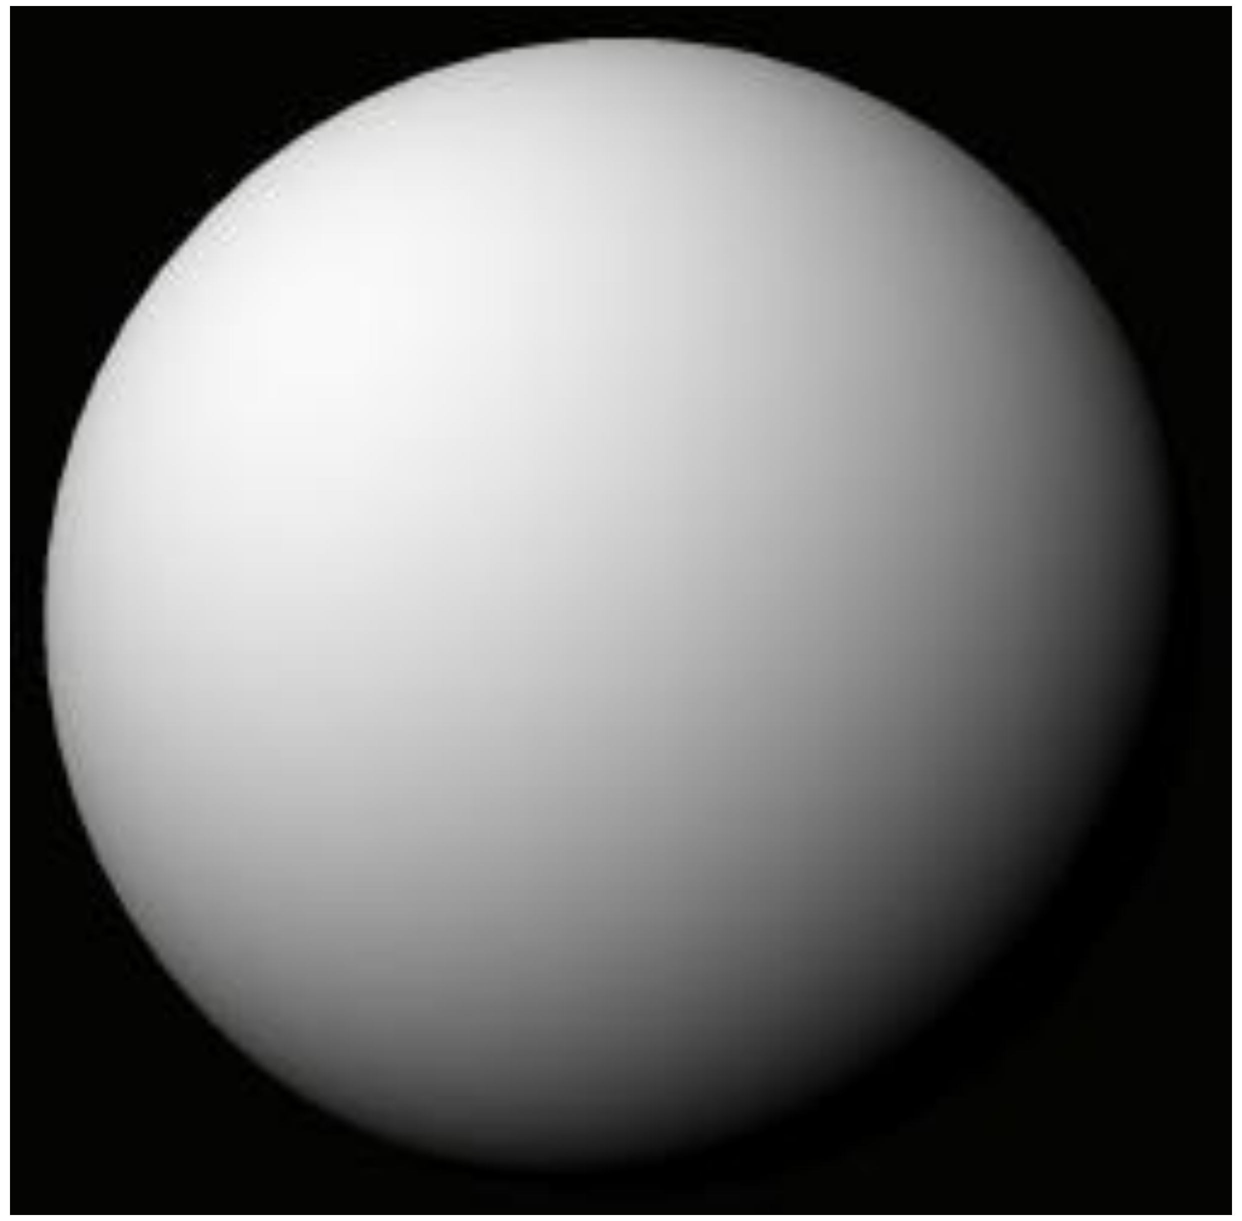

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

Assume k_d is 1 for now.
What can we measure from one image?

- $\cos^{-1}(I)$ is the angle between N and L
- Add assumptions:
  - Constant albedo
  - A few known normals (e.g. silhouettes)
  - Smoothness of normals

In practice, SFS doesn't work very well: assumptions are too restrictive, too much ambiguity in nontrivial scenes.

Source: N. Snavely

An ambiguity that artists exploit!

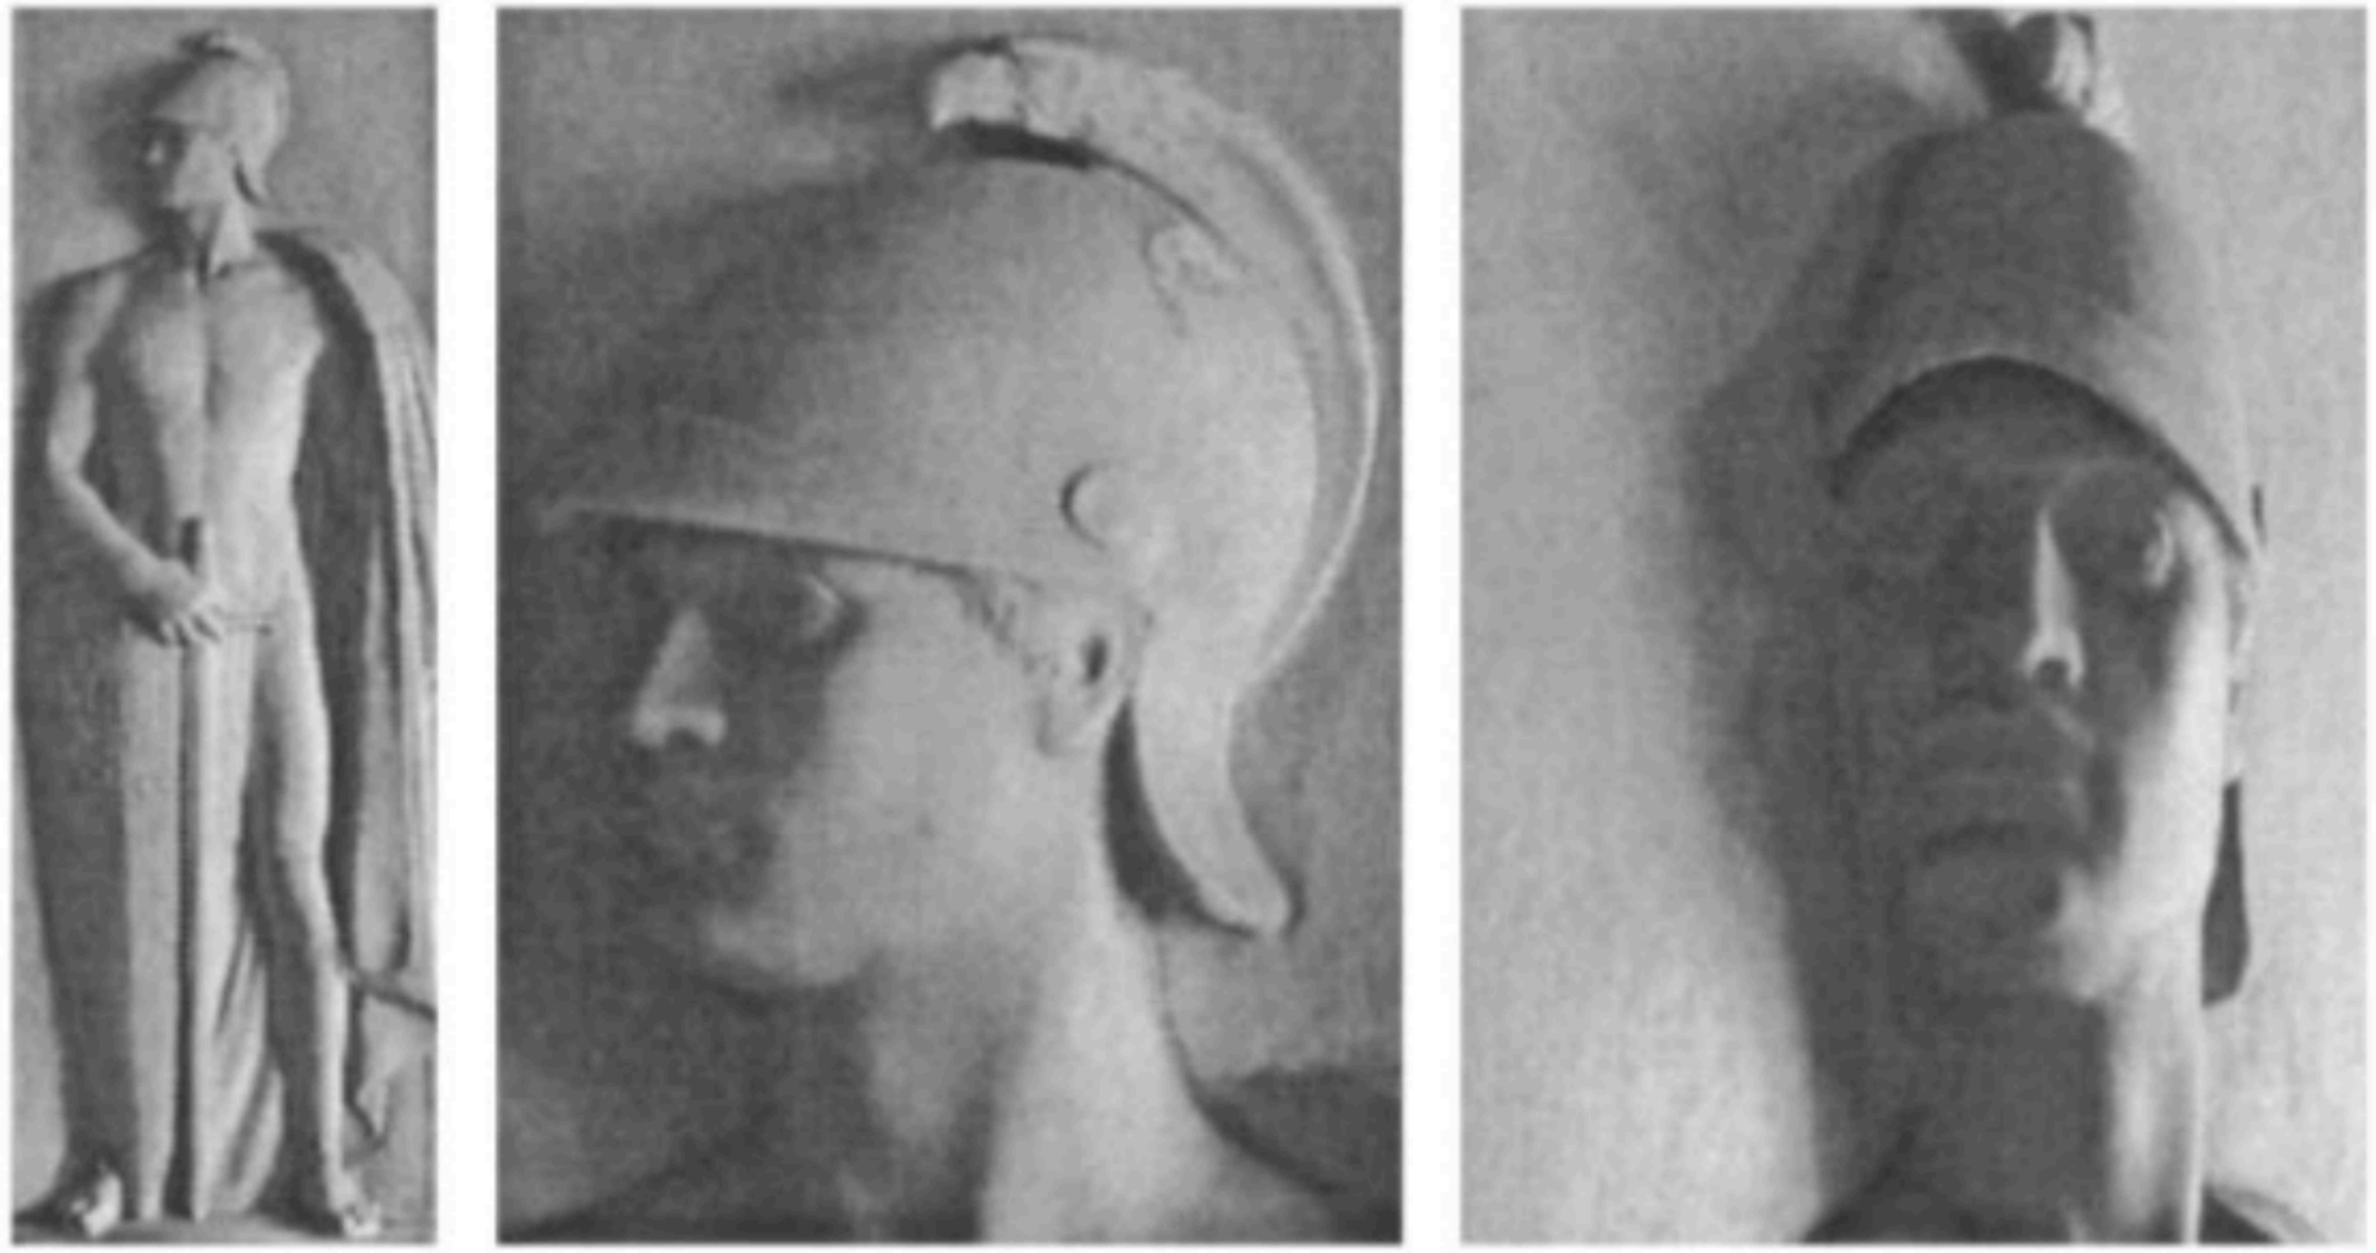

[Belhumeur et al. “The Bas-Relief Ambiguity”, 1999]

Contours provide extra shape information

Consider points on the *occluding contour*:

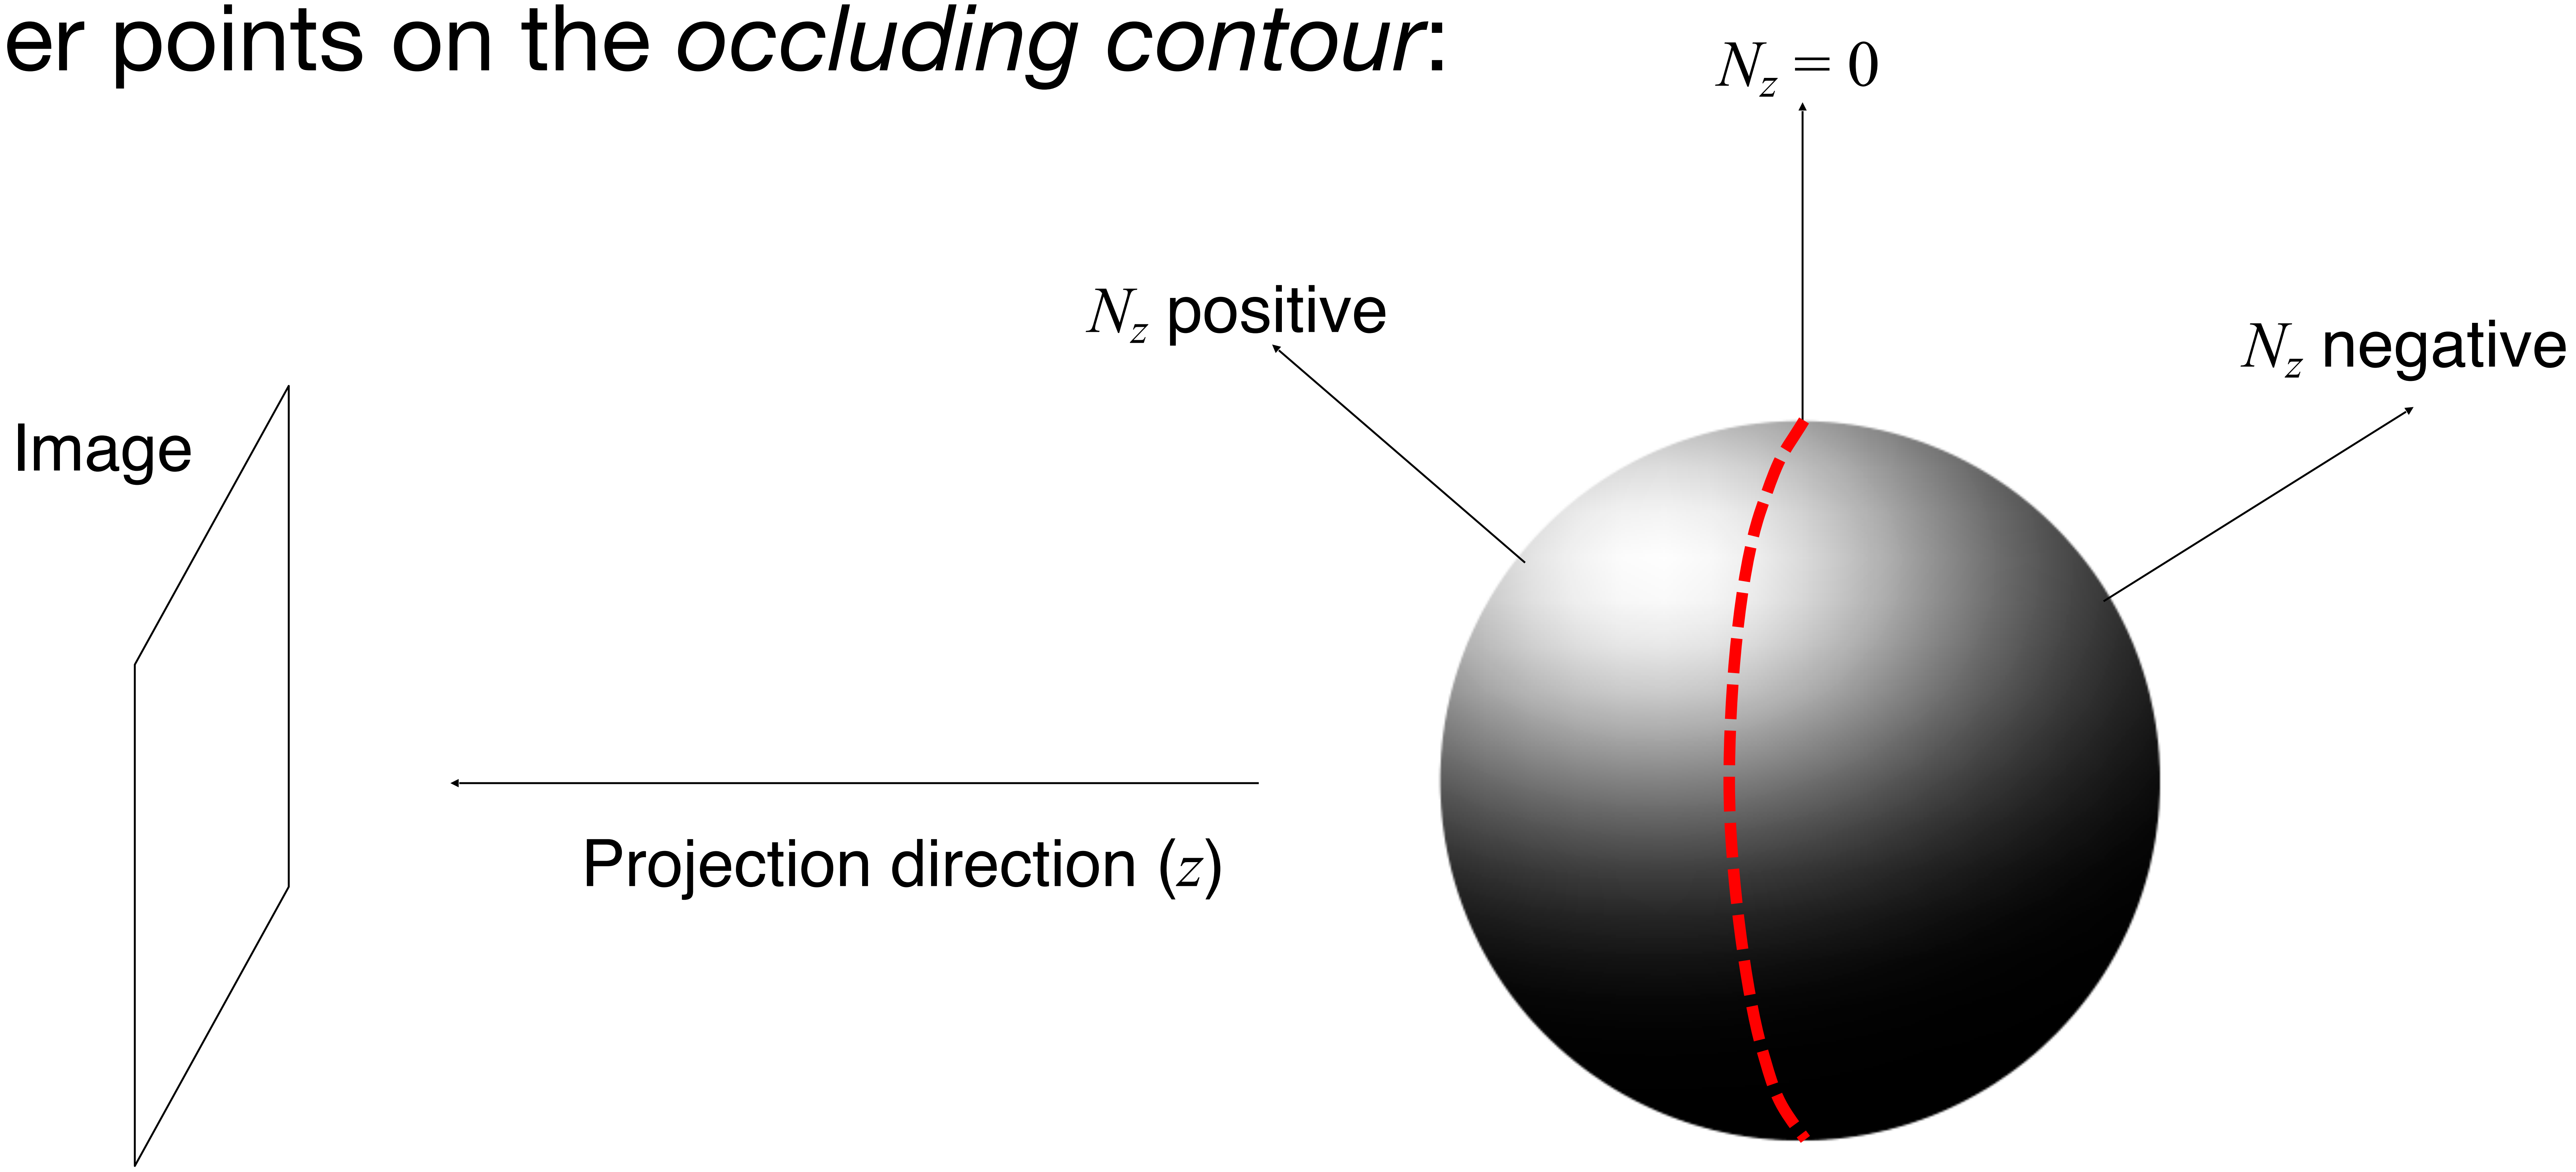

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Source: S. Lazebnik

Application: finding the direction of the light source

$$I(x,y) = \mathbf{N}(x,y) \cdot \mathbf{S}(x,y)$$

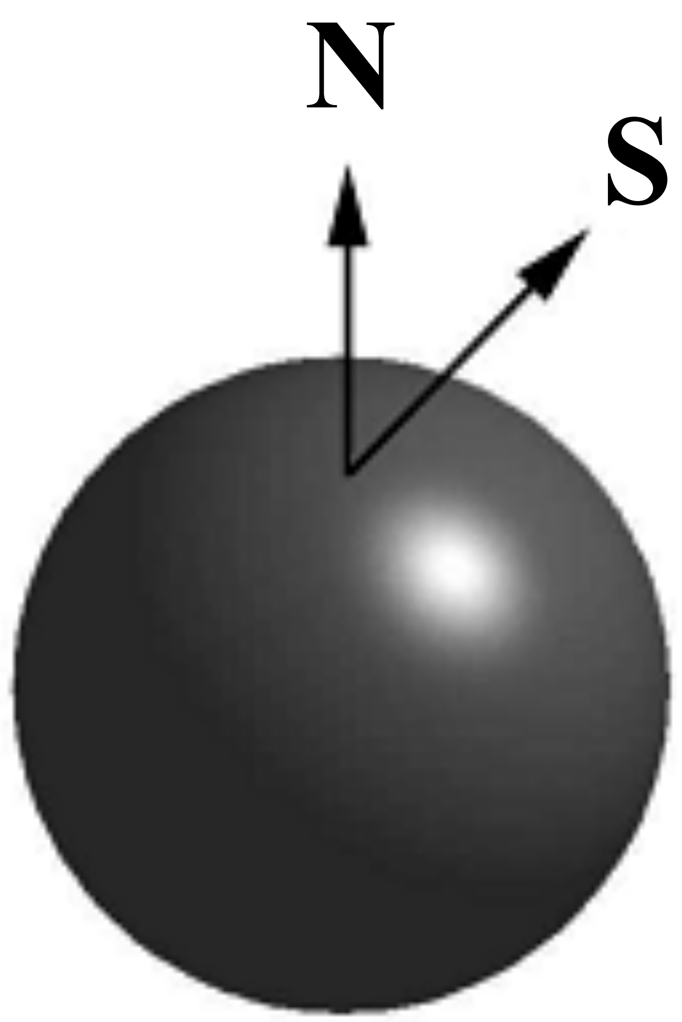

Full 3D case:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) & N_z(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) & N_z(x_2,y_2) \\ \vdots & \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) & N_z(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \\ S_z \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

For points on the occluding contour, $N_z = 0$:

$$\begin{pmatrix} N_x(x_1,y_1) & N_y(x_1,y_1) \\ N_x(x_2,y_2) & N_y(x_2,y_2) \\ \vdots & \vdots \\ N_x(x_n,y_n) & N_y(x_n,y_n) \end{pmatrix} \begin{pmatrix} S_x \\ S_y \end{pmatrix} = \begin{pmatrix} I(x_1,y_1) \\ I(x_2,y_2) \\ \vdots \\ I(x_n,y_n) \end{pmatrix}$$

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Source: S. Lazebnik

Finding the direction of the light source

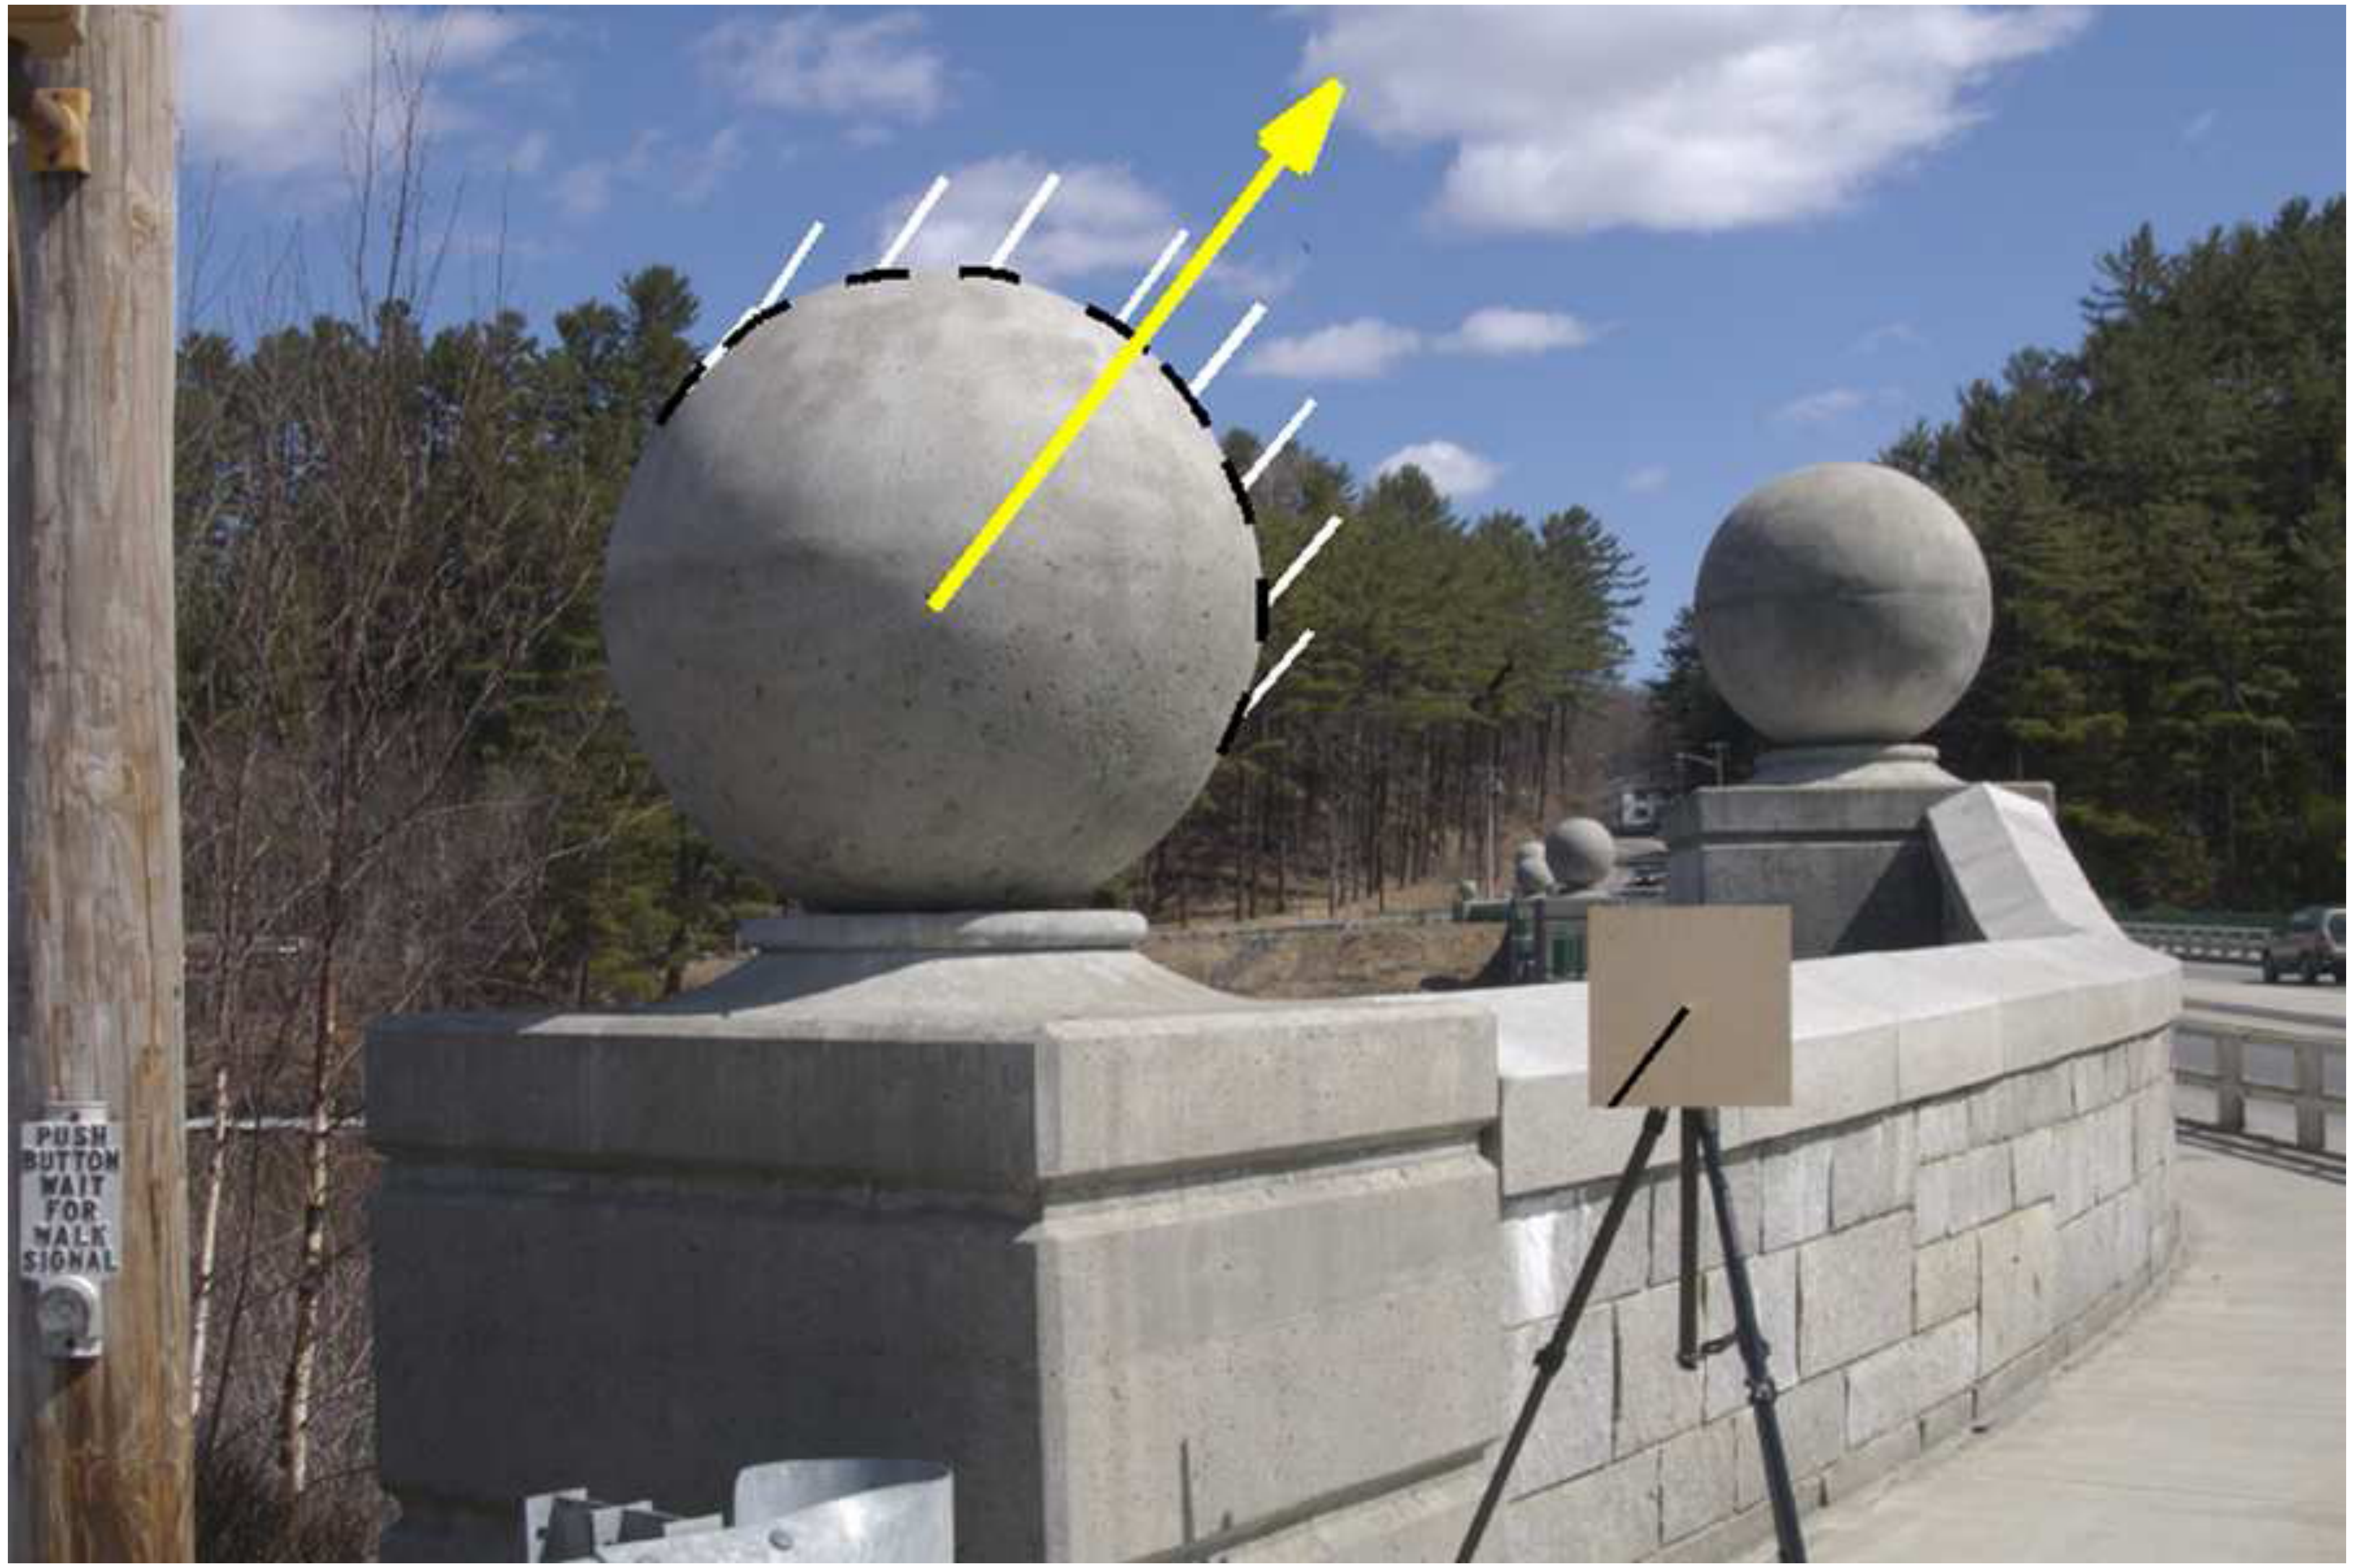

P. Nillius and J.-O. Eklundh, “Automatic estimation of the projected light source direction,” CVPR 2001

Source: S. Lazebnik

Application: Detecting composite photos

Fake photo

Real photo

[Johnson and Farid, 2005]

Source: S. Lazebnik

Photometric stereo

Source: N. Snavely

Photometric stereo

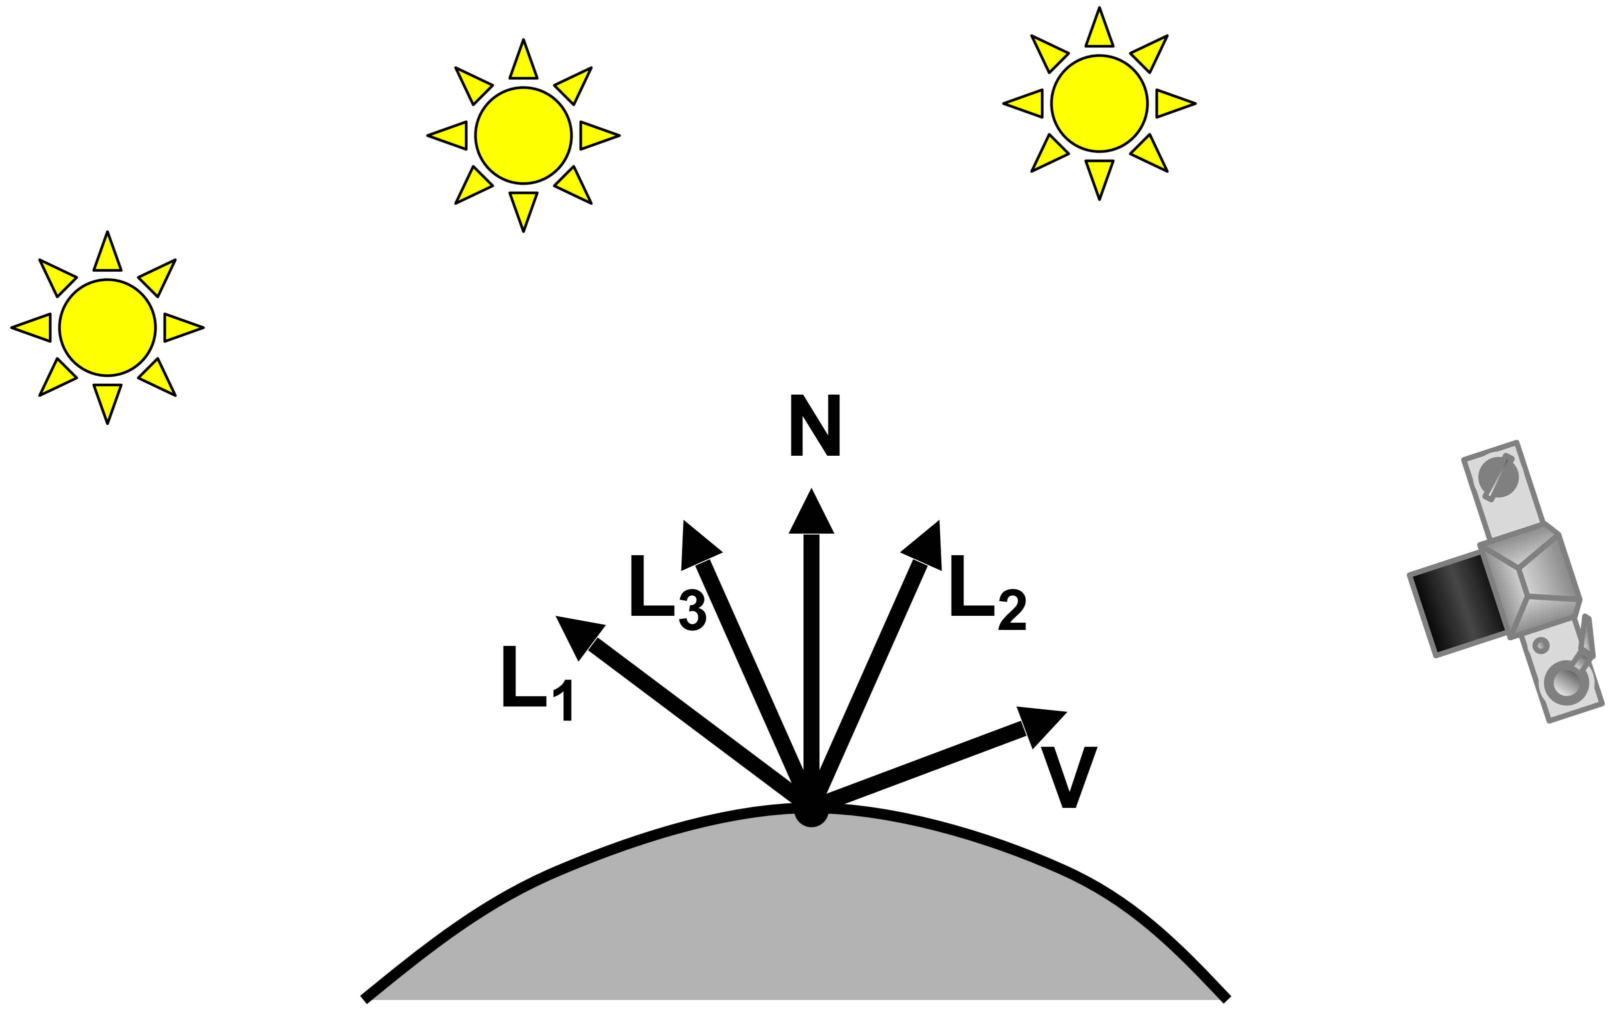

$$
\begin{aligned}
I_1 &= k_d \mathbf{N} \cdot \mathbf{L_1} \\
I_2 &= k_d \mathbf{N} \cdot \mathbf{L_2} \\
I_3 &= k_d \mathbf{N} \cdot \mathbf{L_3}
\end{aligned}
$$

Can write this as a linear system, and solve:

$$
\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = k_d \begin{bmatrix} \mathbf{L_1}^T \\ \mathbf{L_2}^T \\ \mathbf{L_3}^T \end{bmatrix} \mathbf{N}
$$

Source: N. Snavely

Photometric Stereo

Input

Recovered albedo

Recovered normal field

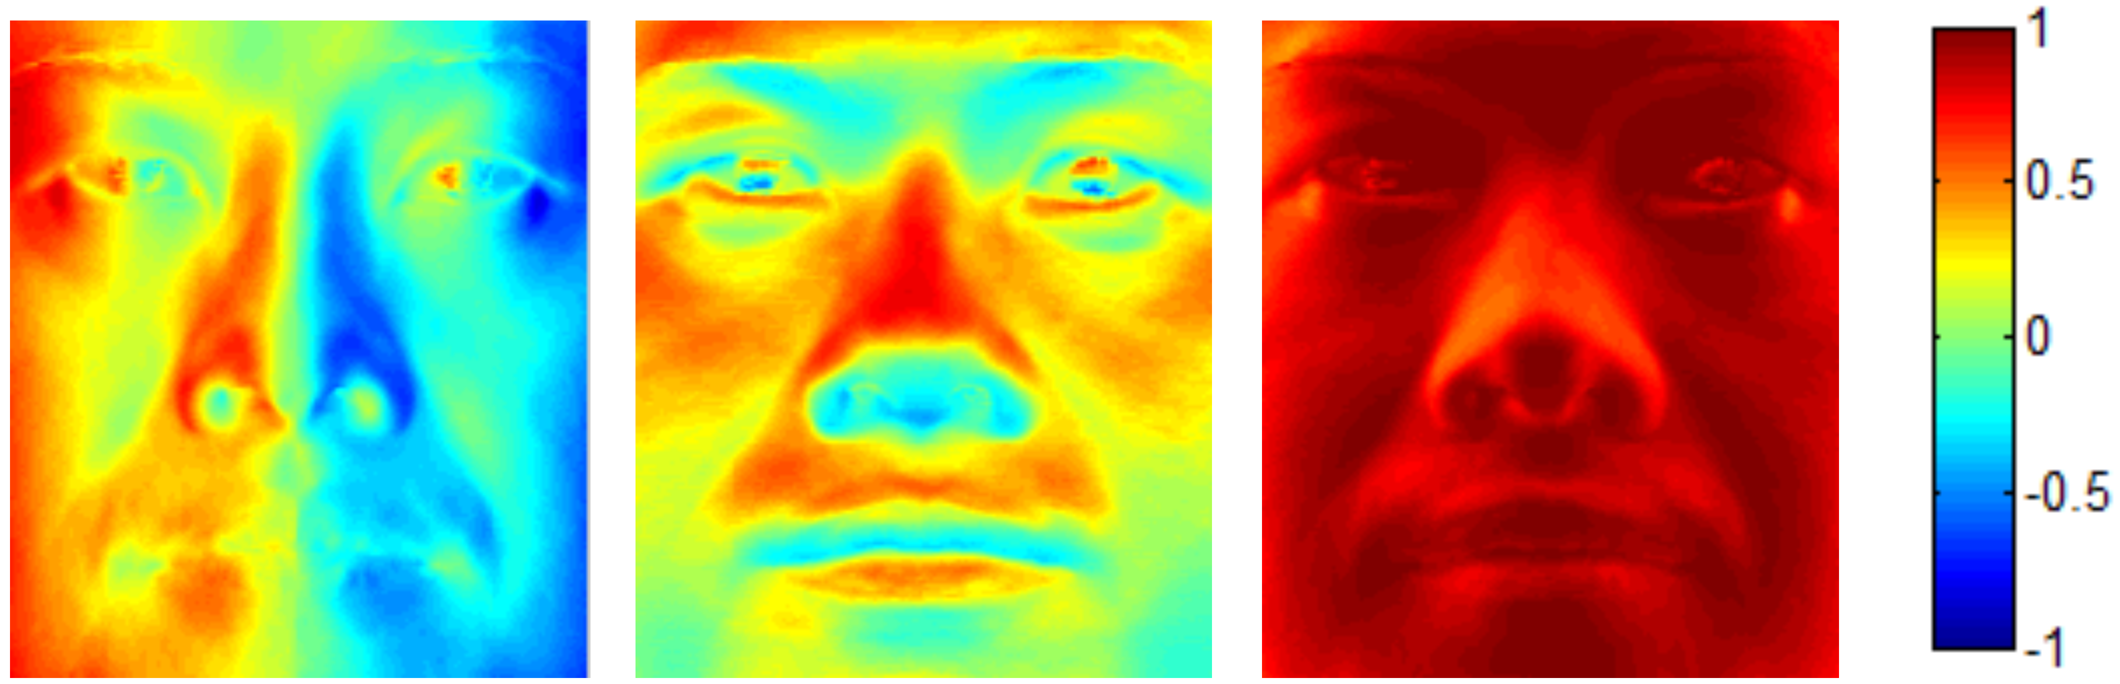

x y z

Recovered surface model

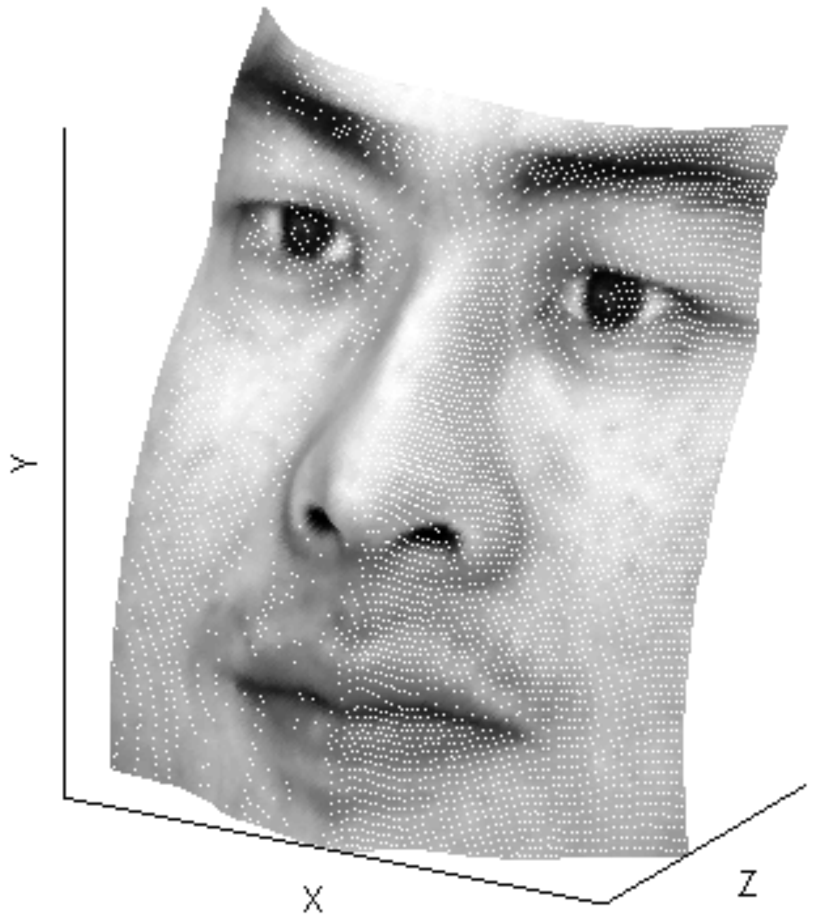

Source: Forsyth & Ponce, S. Lazebnik

Photometric Stereo

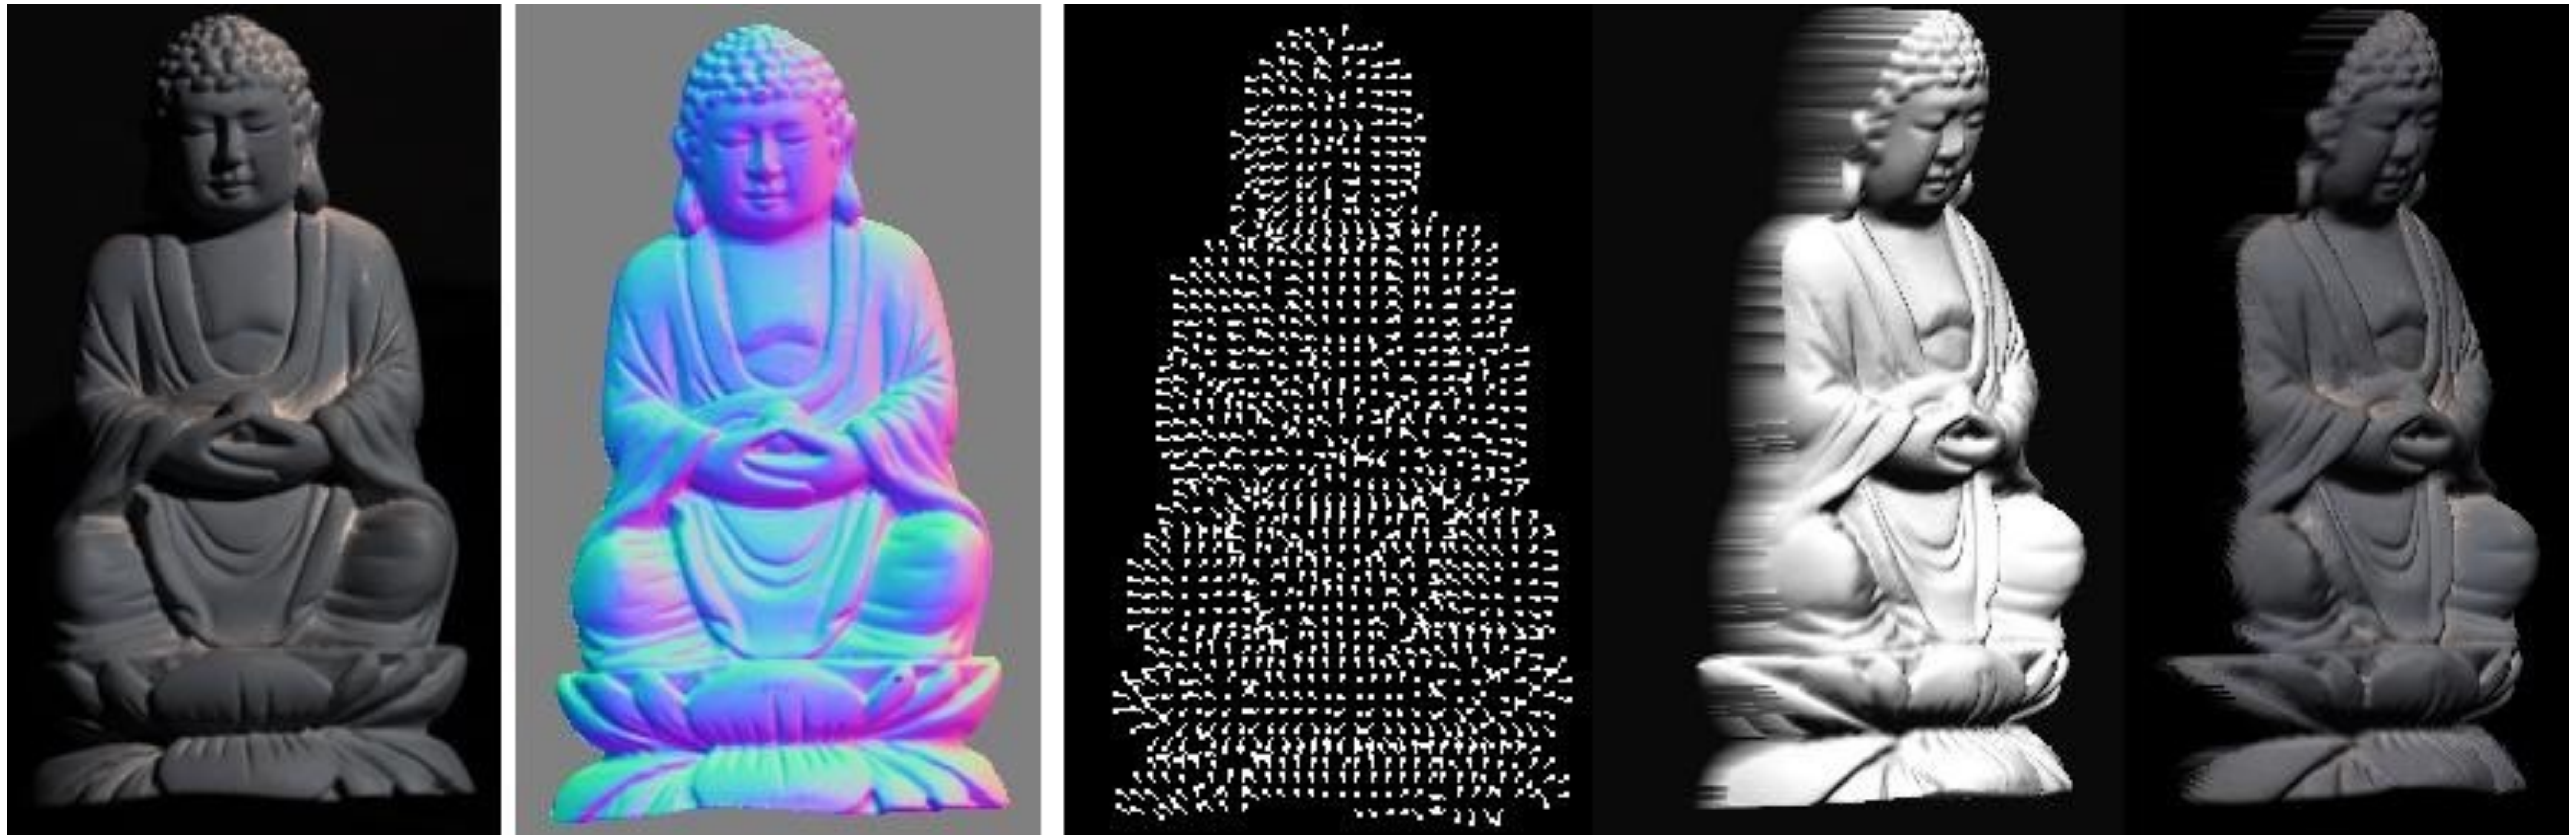

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Source: N. Snavely

Video photometric stereo

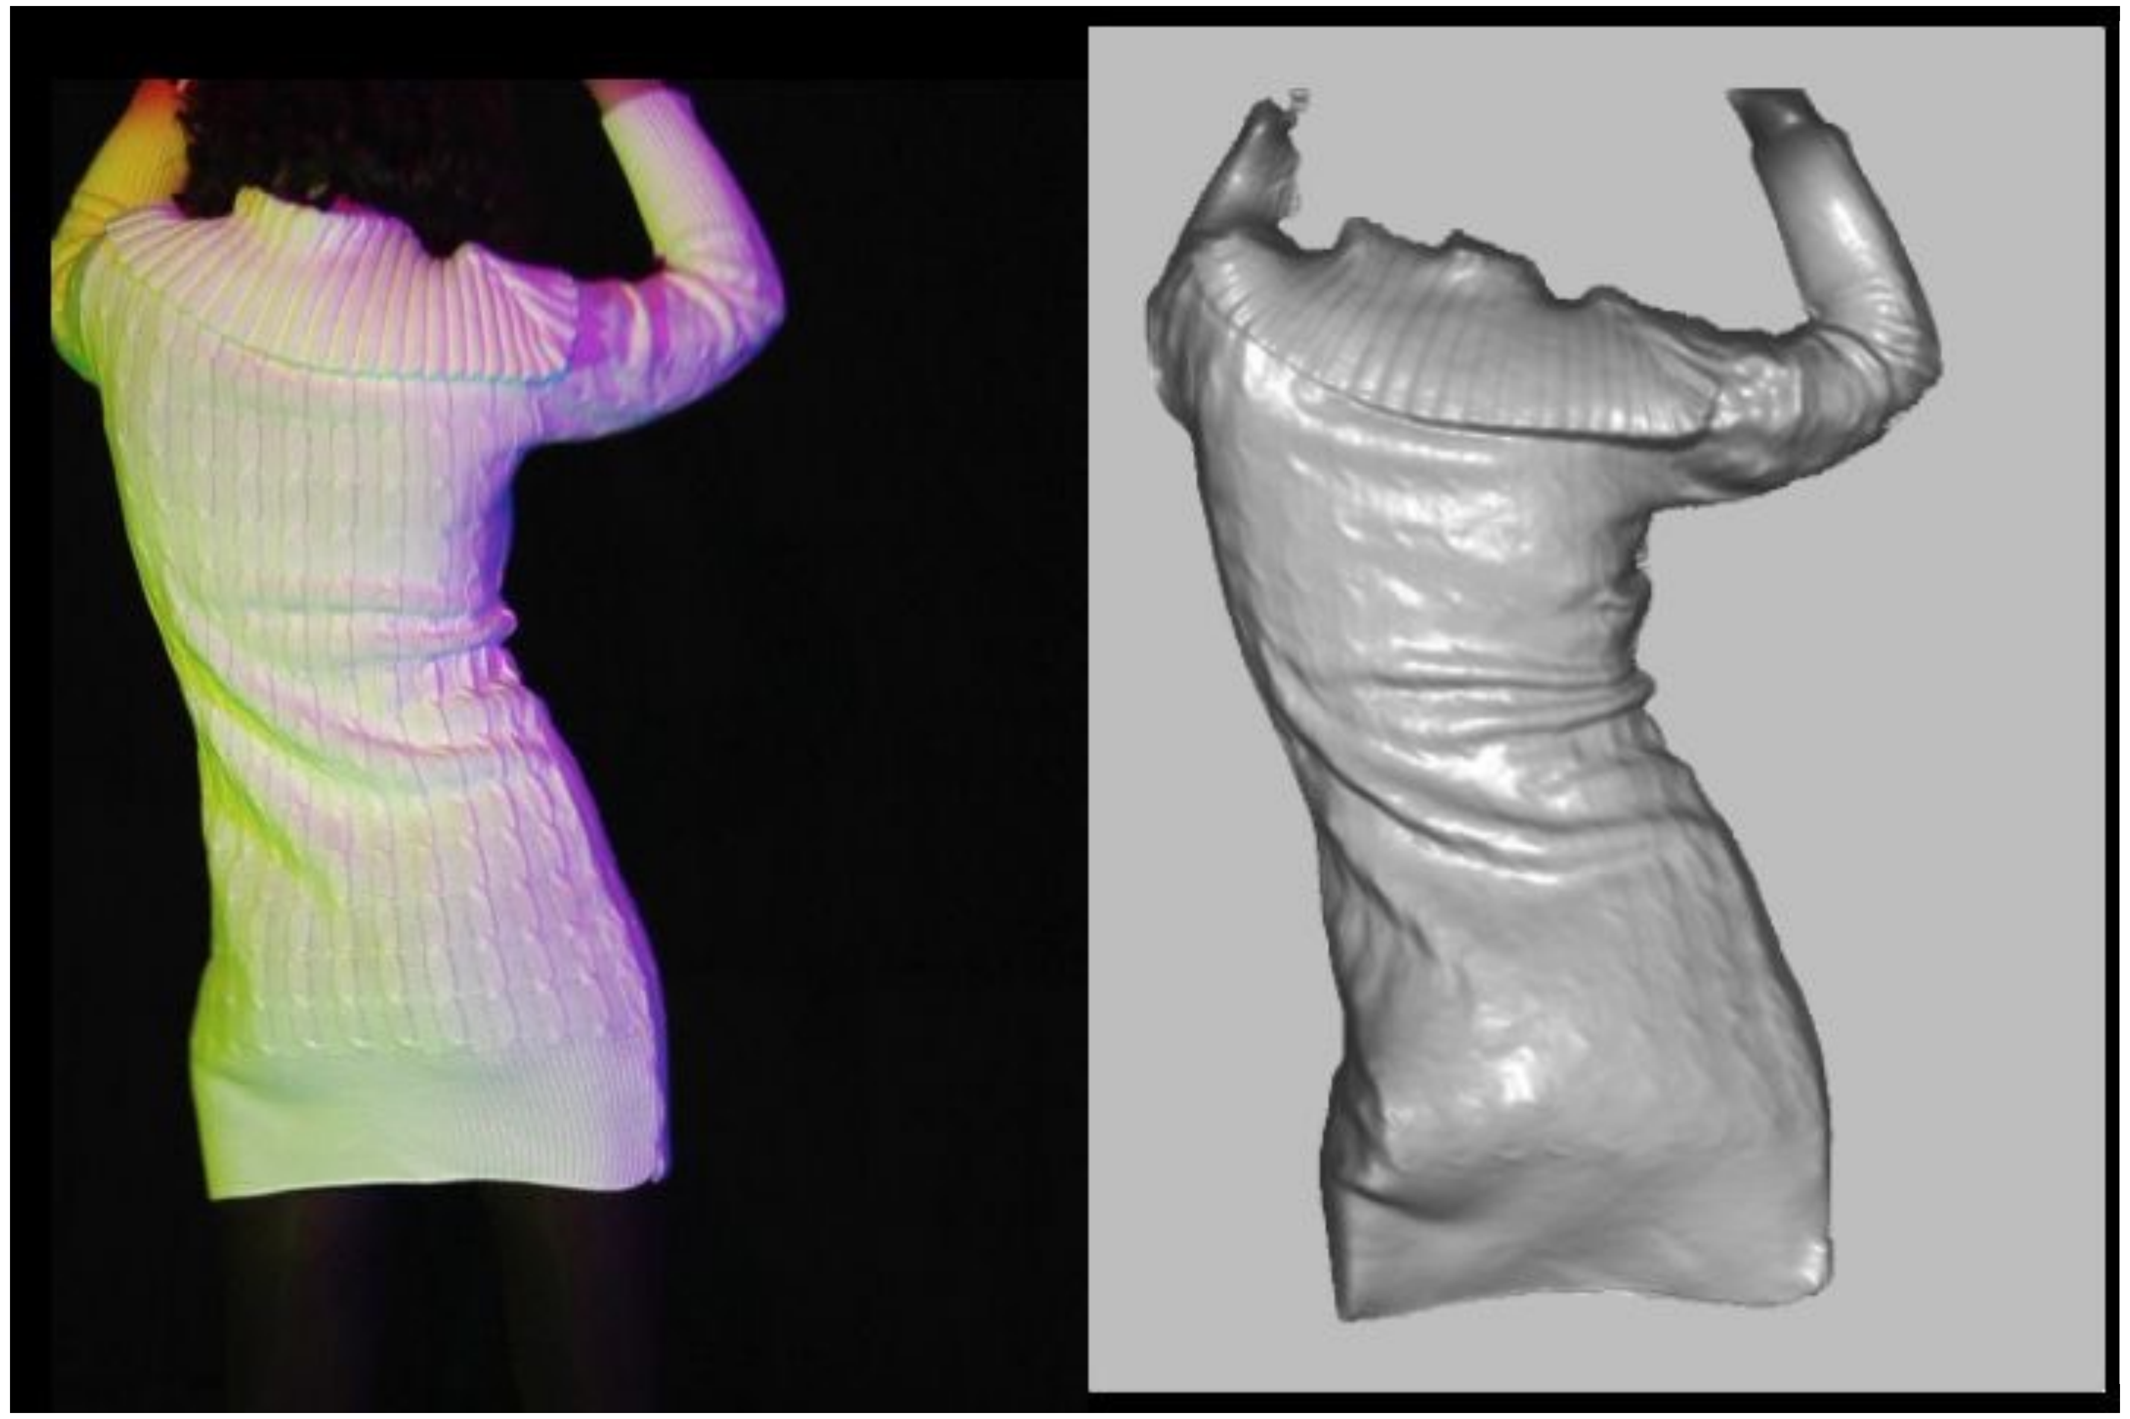

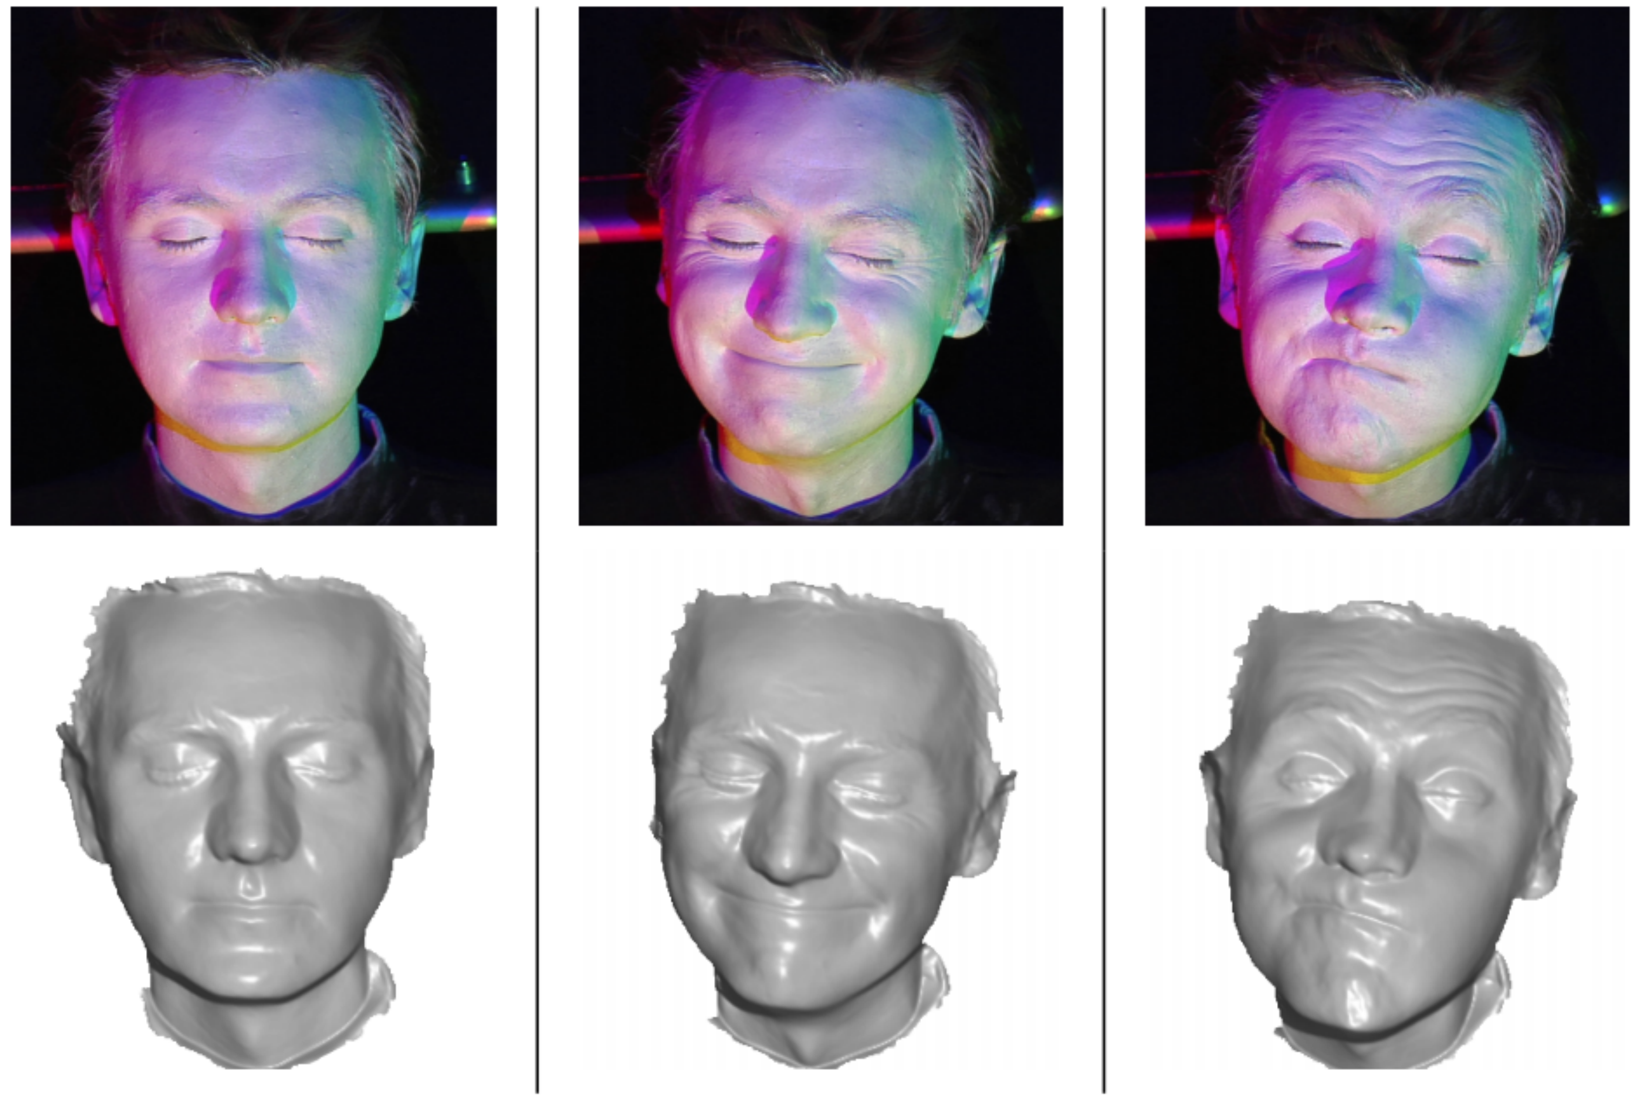

Fig. 2. **Applying the original algorithm to a face with white makeup.** Top: example input frames from video of an actor smiling and grimacing. Bottom: the resulting integrated surfaces.

Video Normals from Colored Lights
Gabriel J. Brostow, Carlos Hernández, George Vogiatzis, Björn Stenger, Roberto Cipolla
IEEE TPAMI, Vol. 33, No. 10, pages 2104-2114, October 2011.

Source: N. Snavely

But what if we don't know the BRDF?

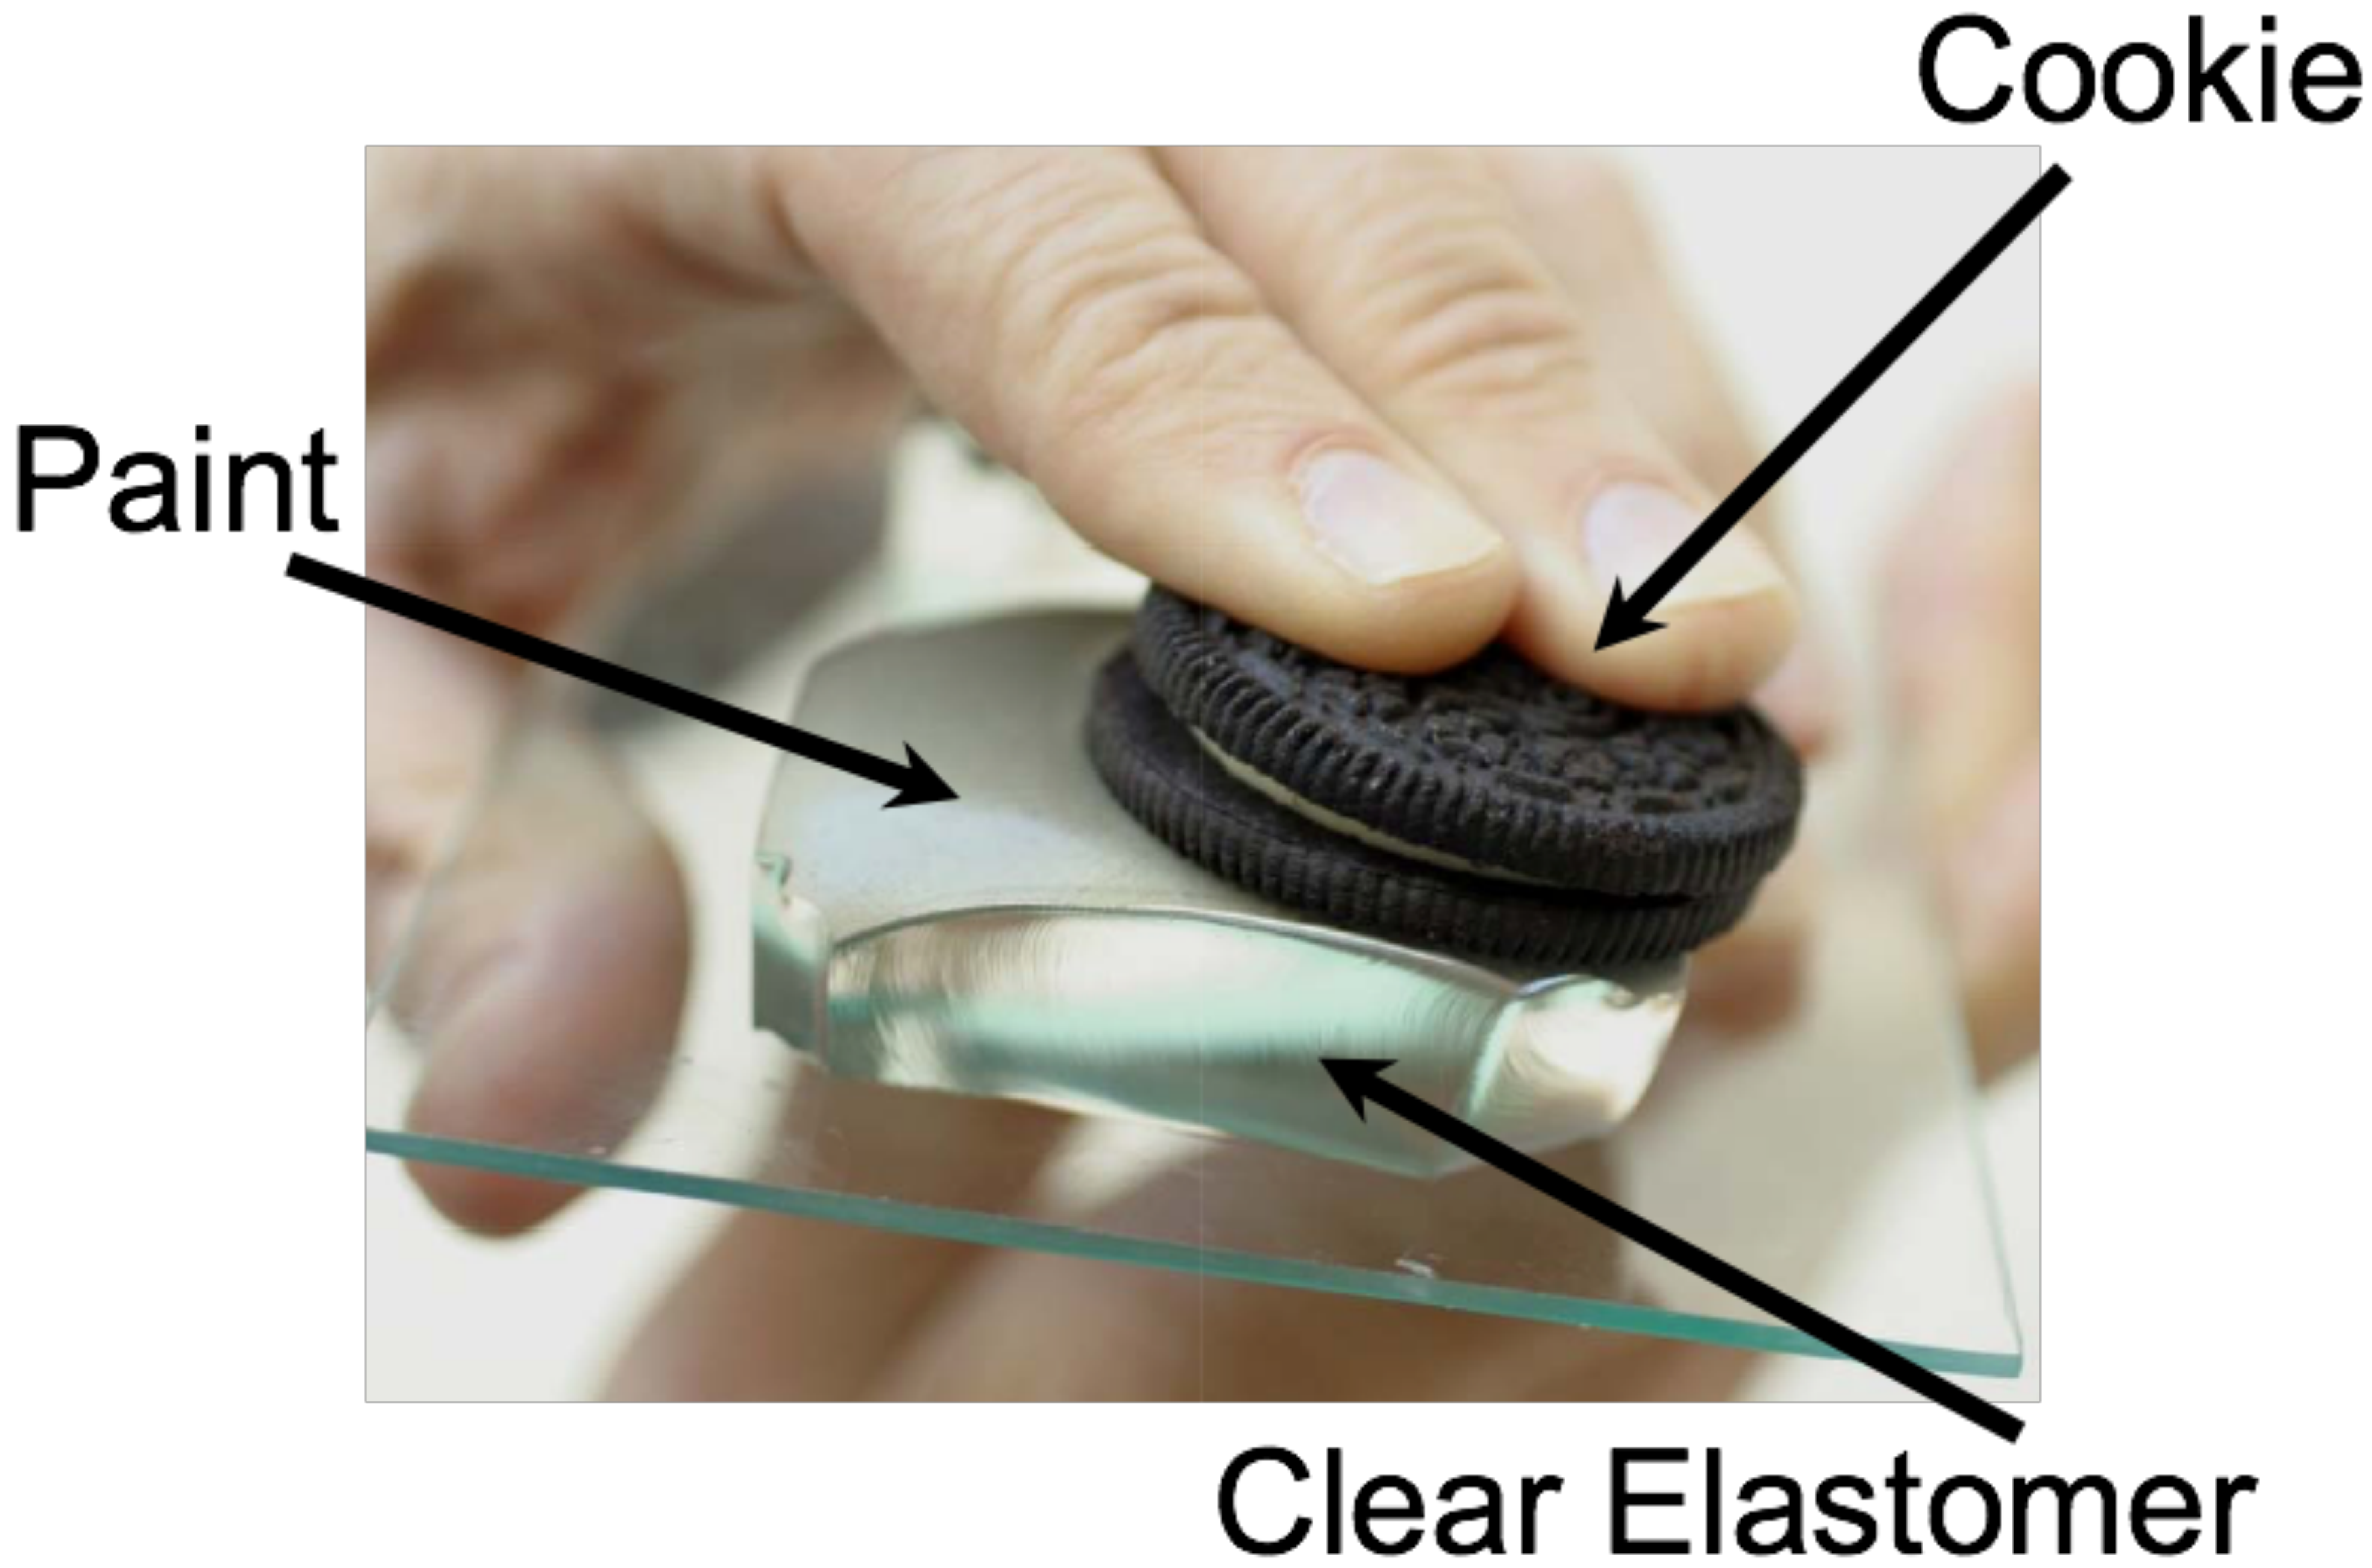

[Johnson and Adelson, 2009]

Source: N. Snavely


OREO


OREO

Lights, camera, action

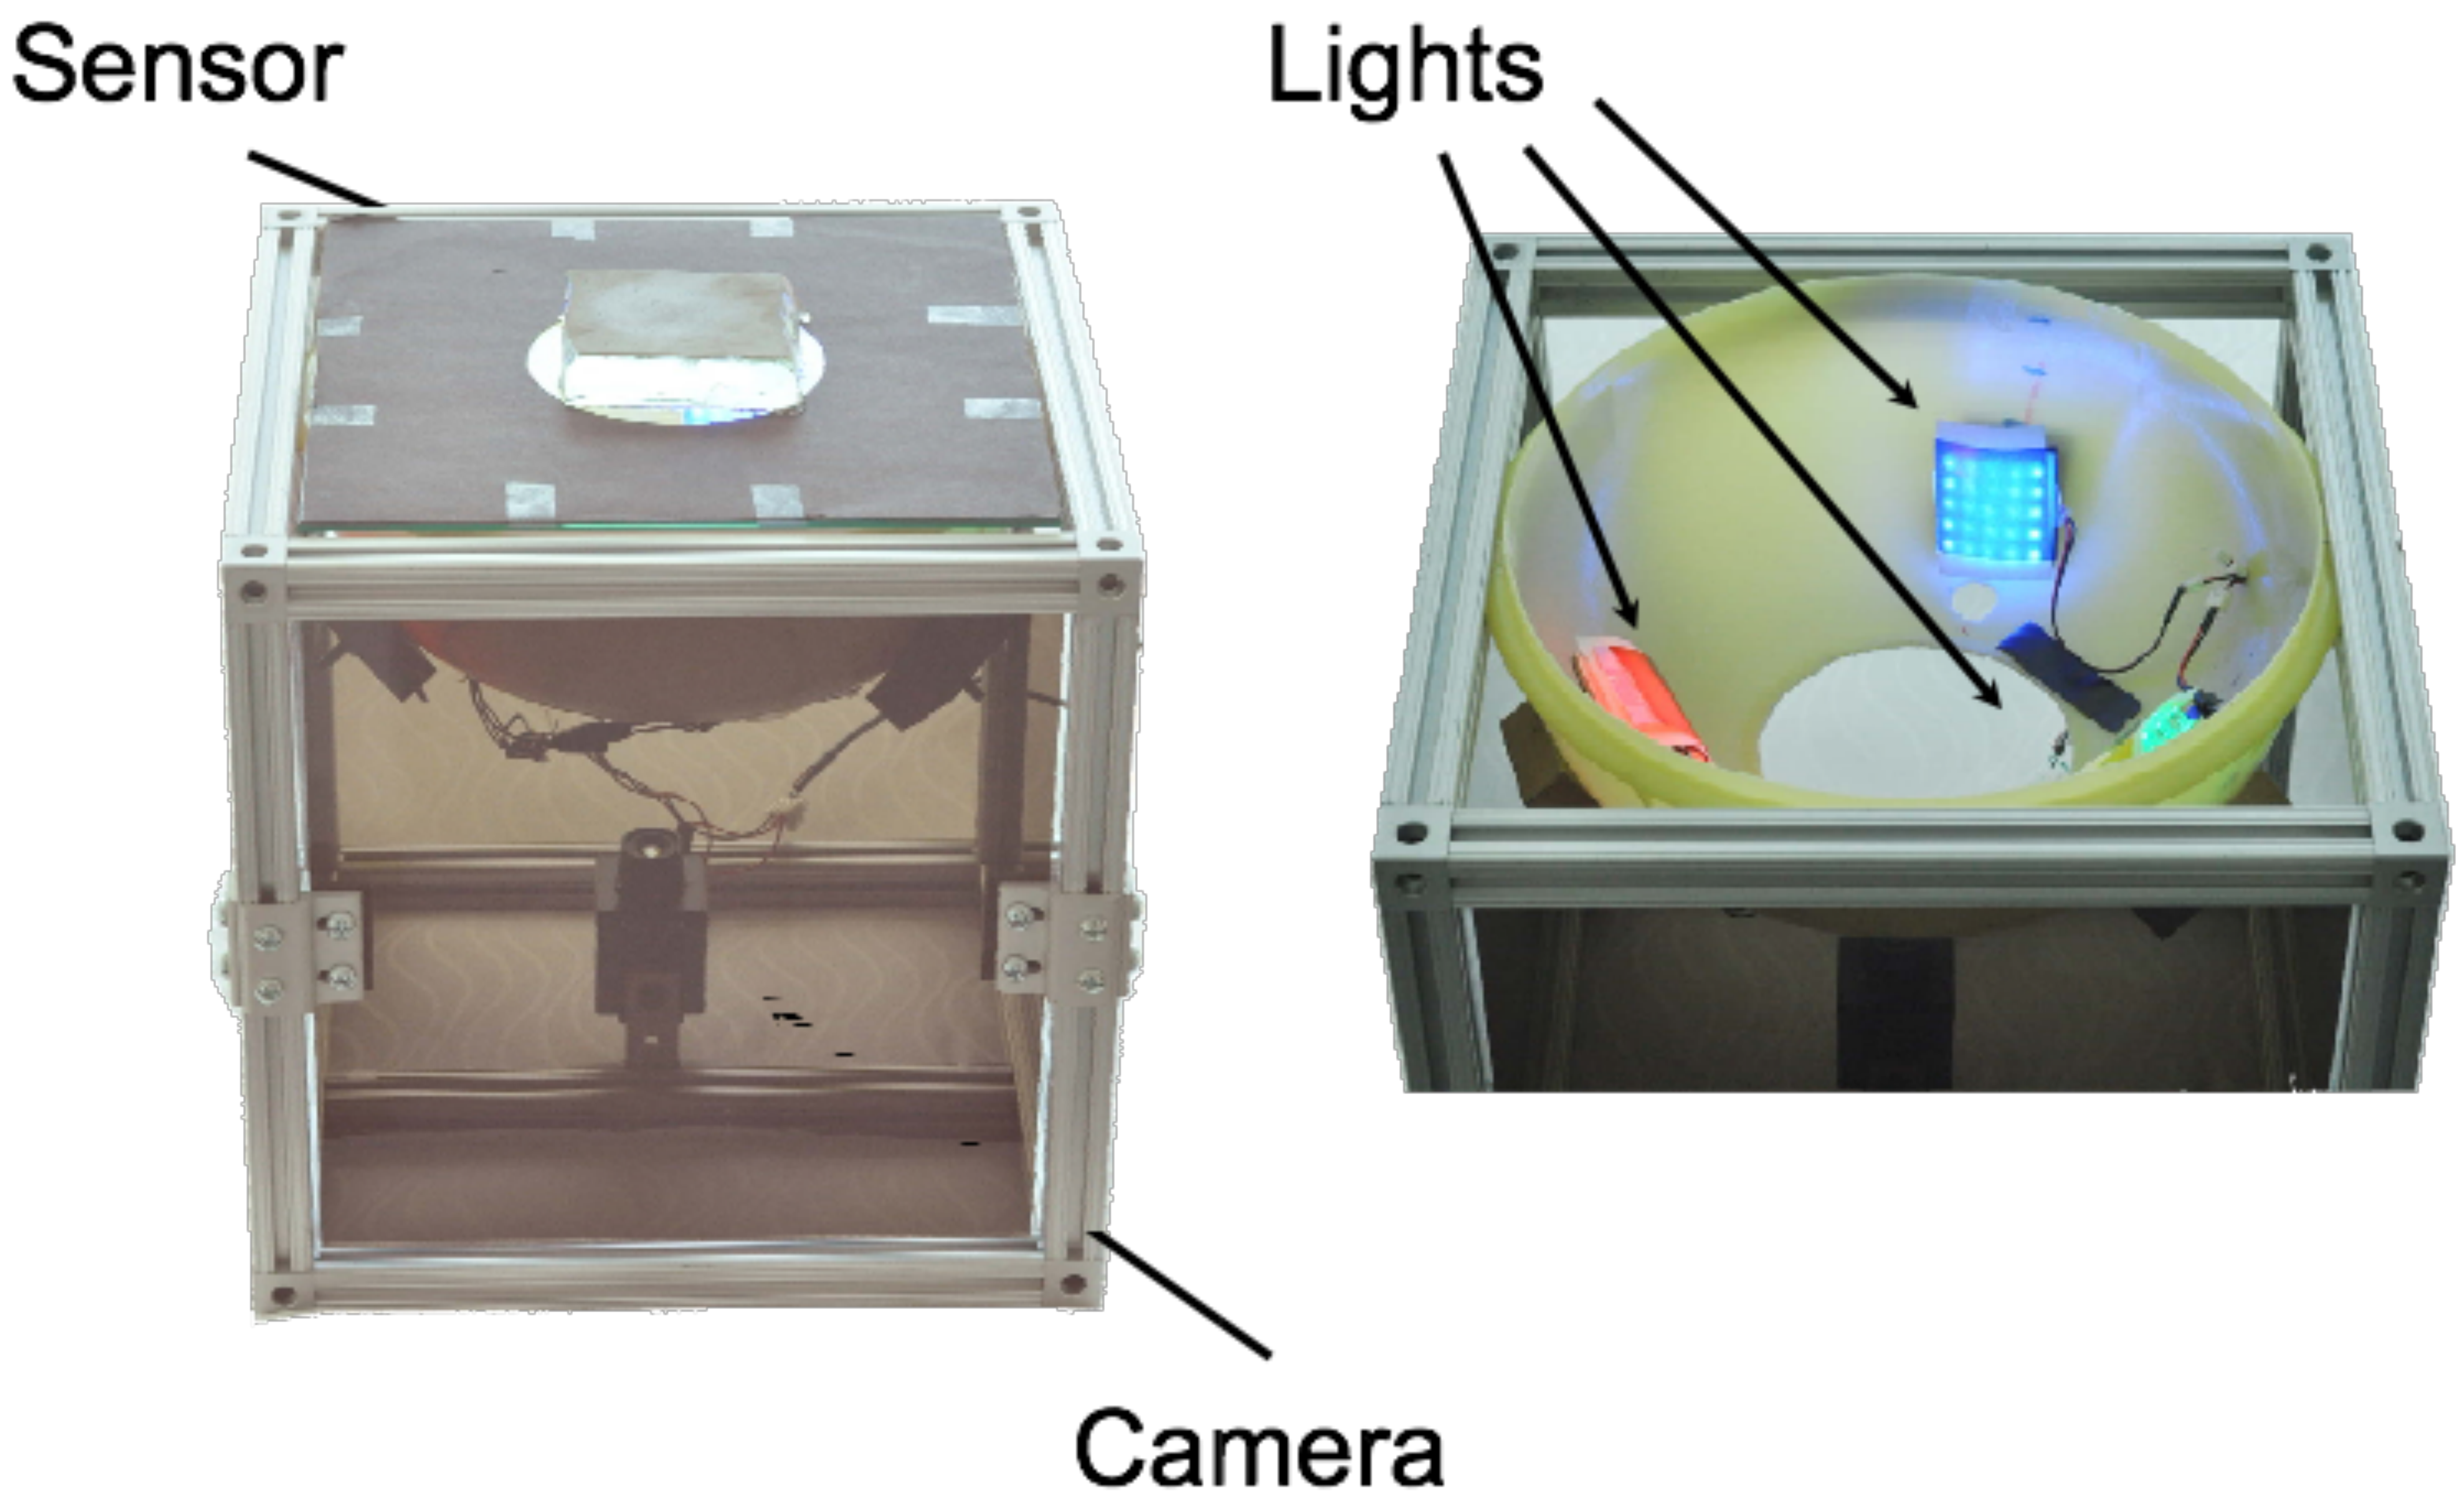

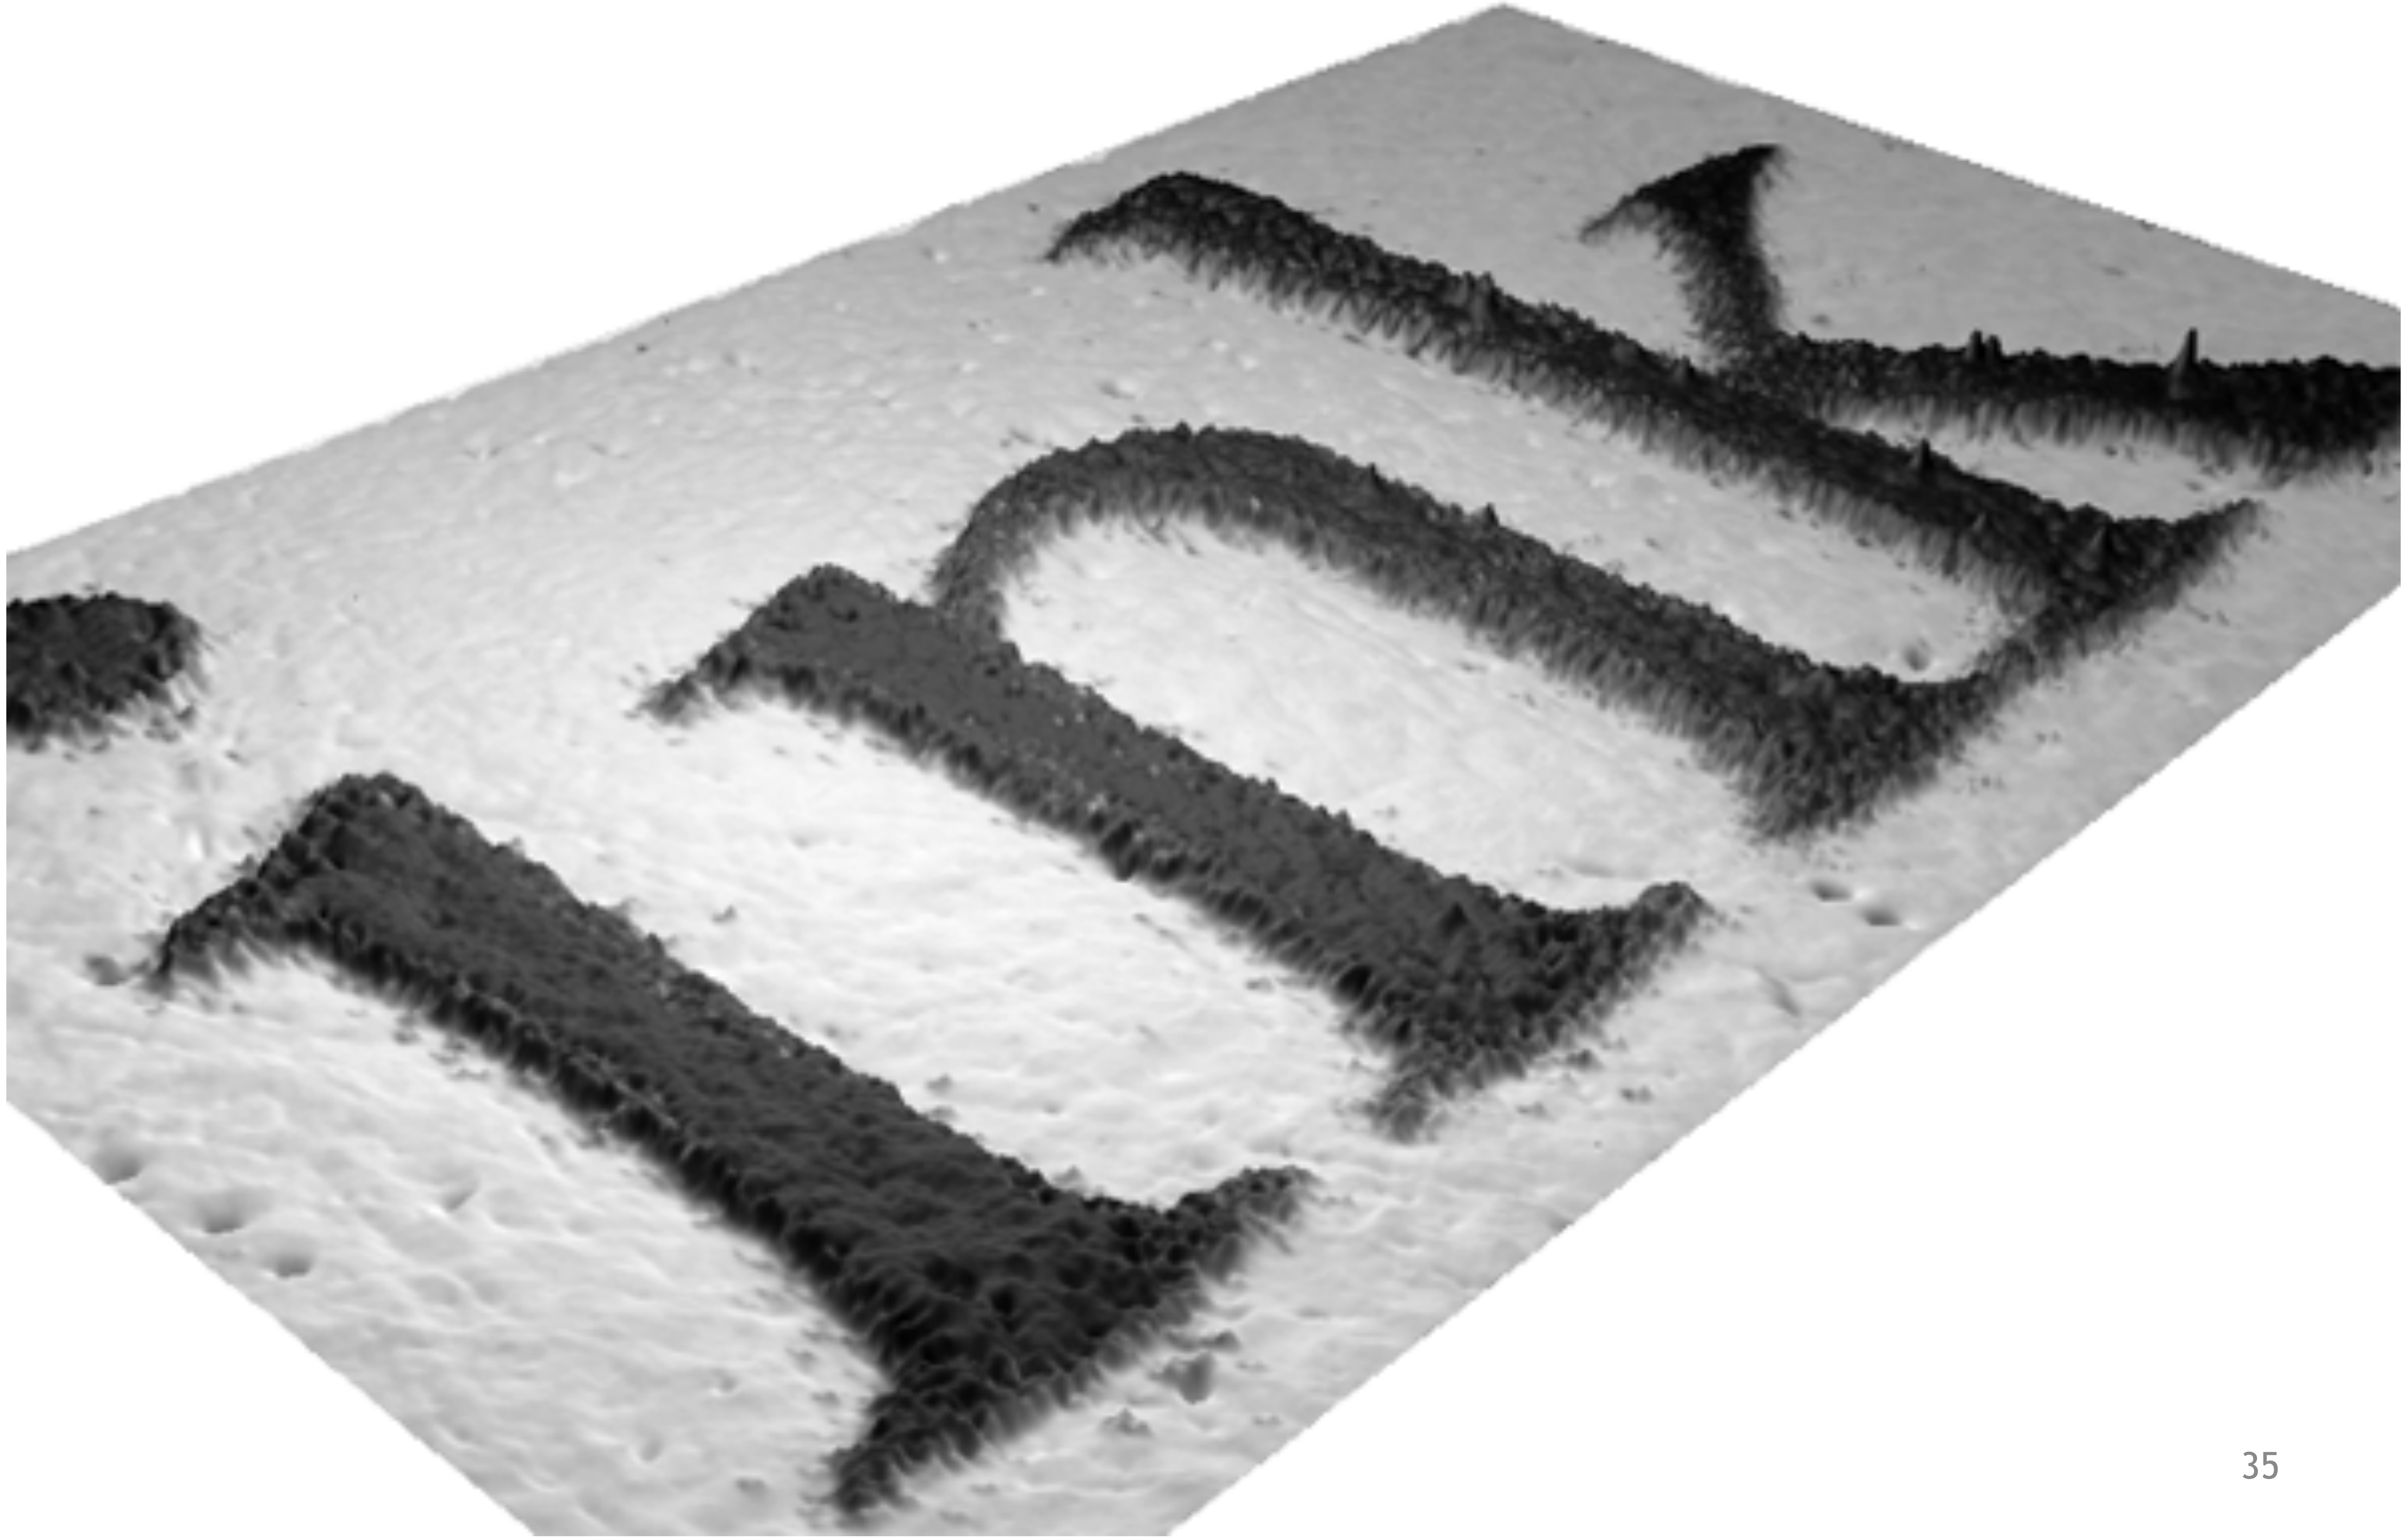

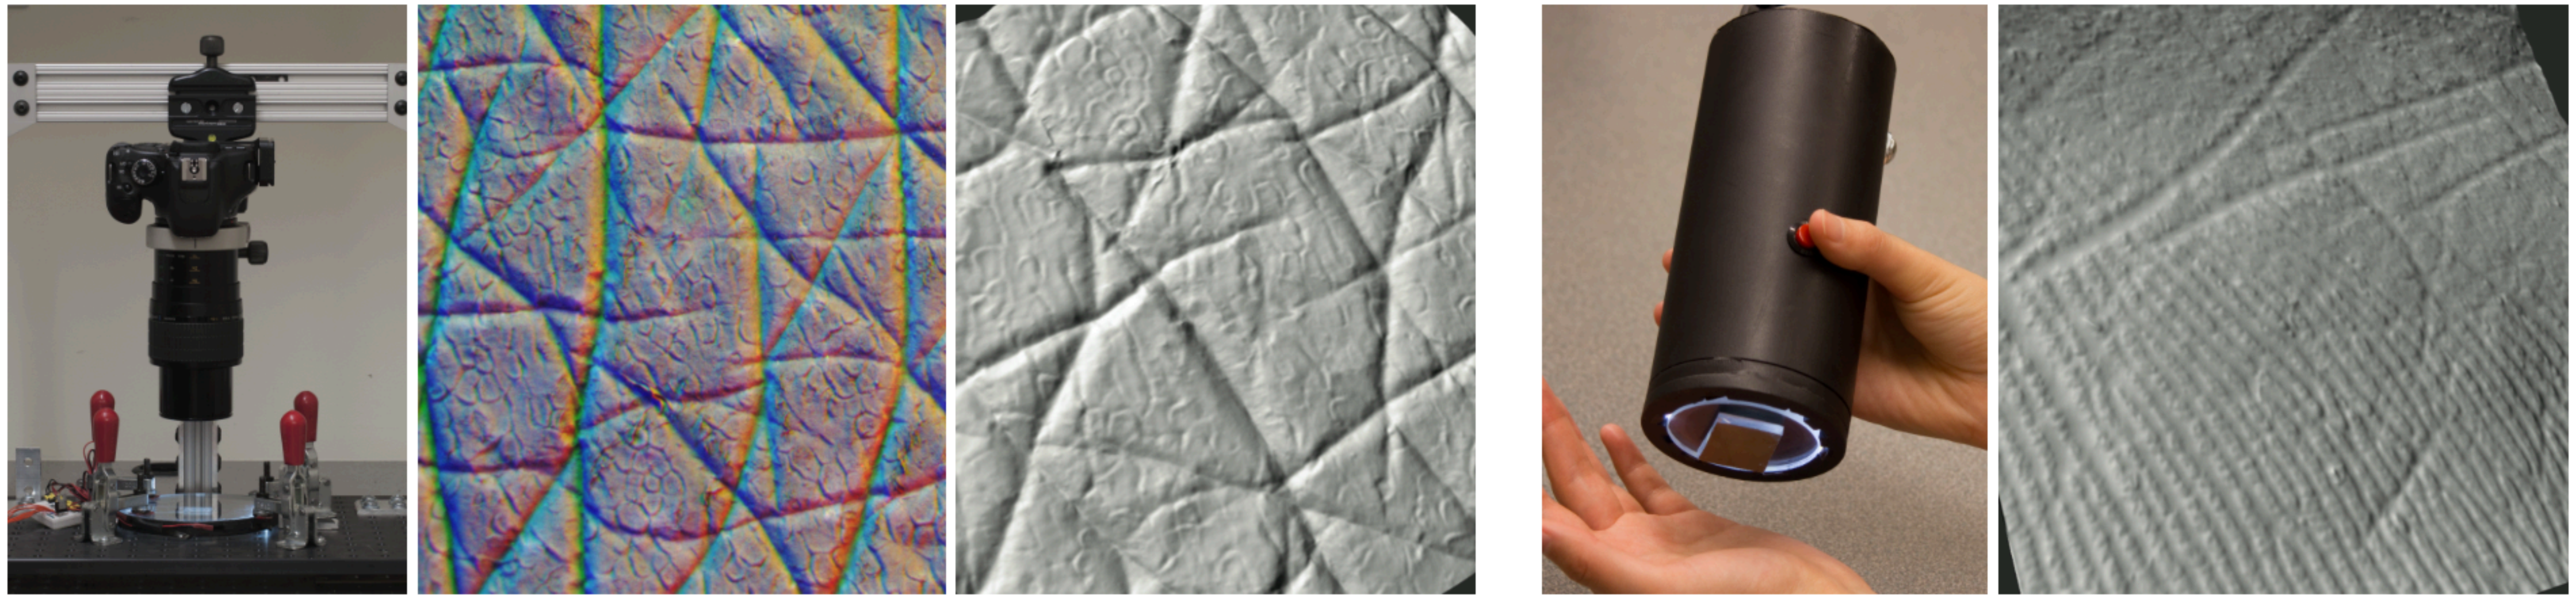

(a) bench configuration (b) captured (c) reconstruction (d) portable configuration (e) reconstruction

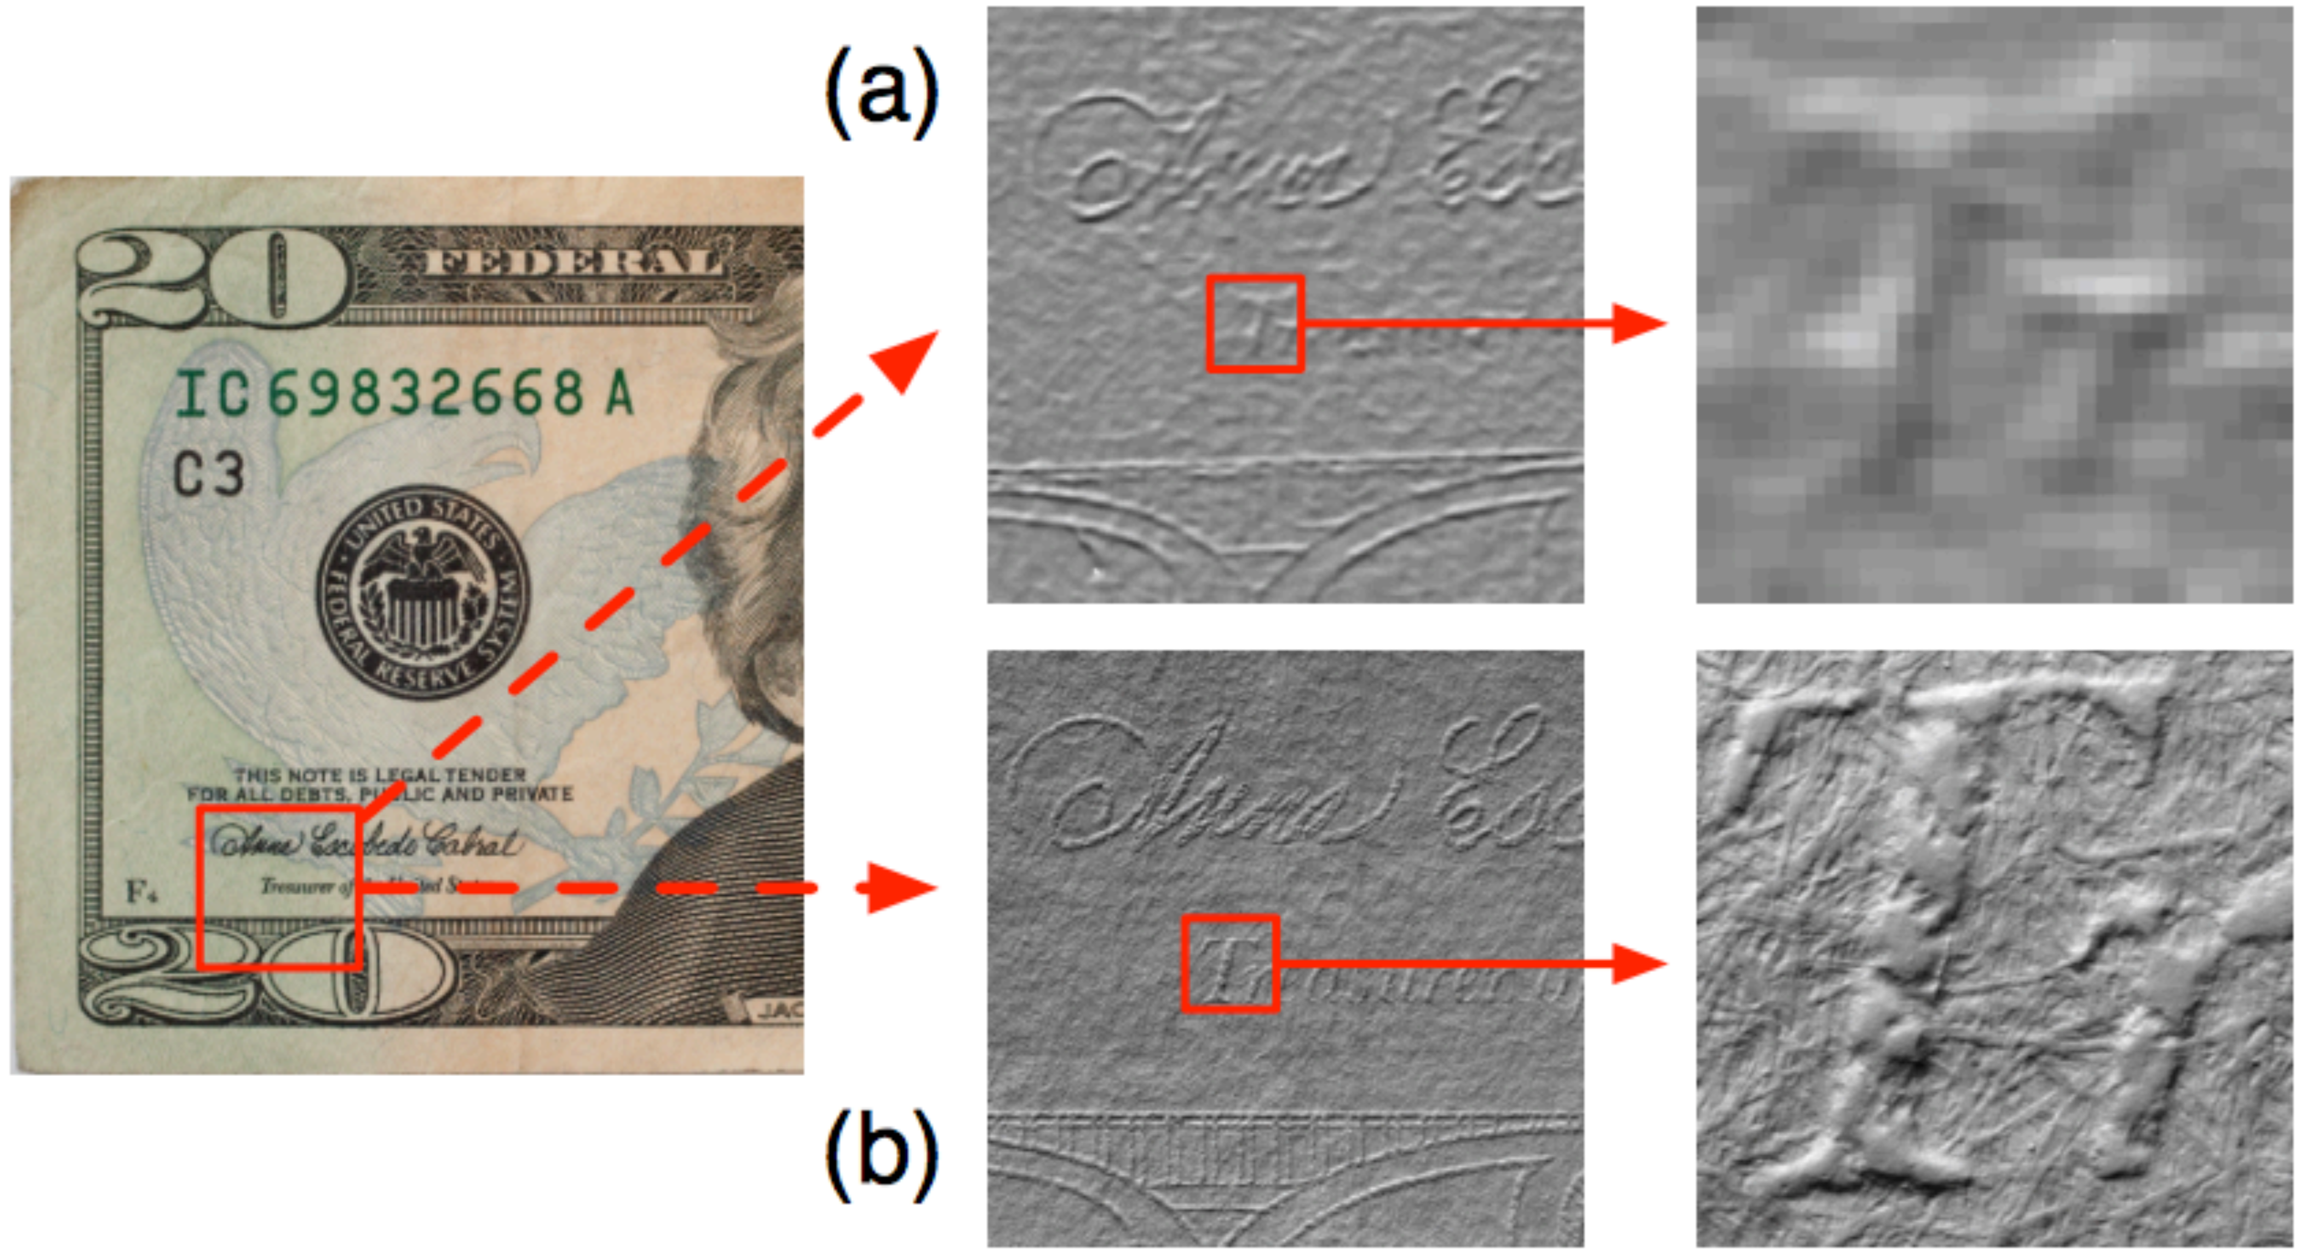

Figure 7: *Comparison with the high-resolution result from the original retrographic sensor. (a) Rendering of the high-resolution \$20 bill example from the original retrographic sensor with a close-up view. (b) Rendering of the captured geometry using our method.*

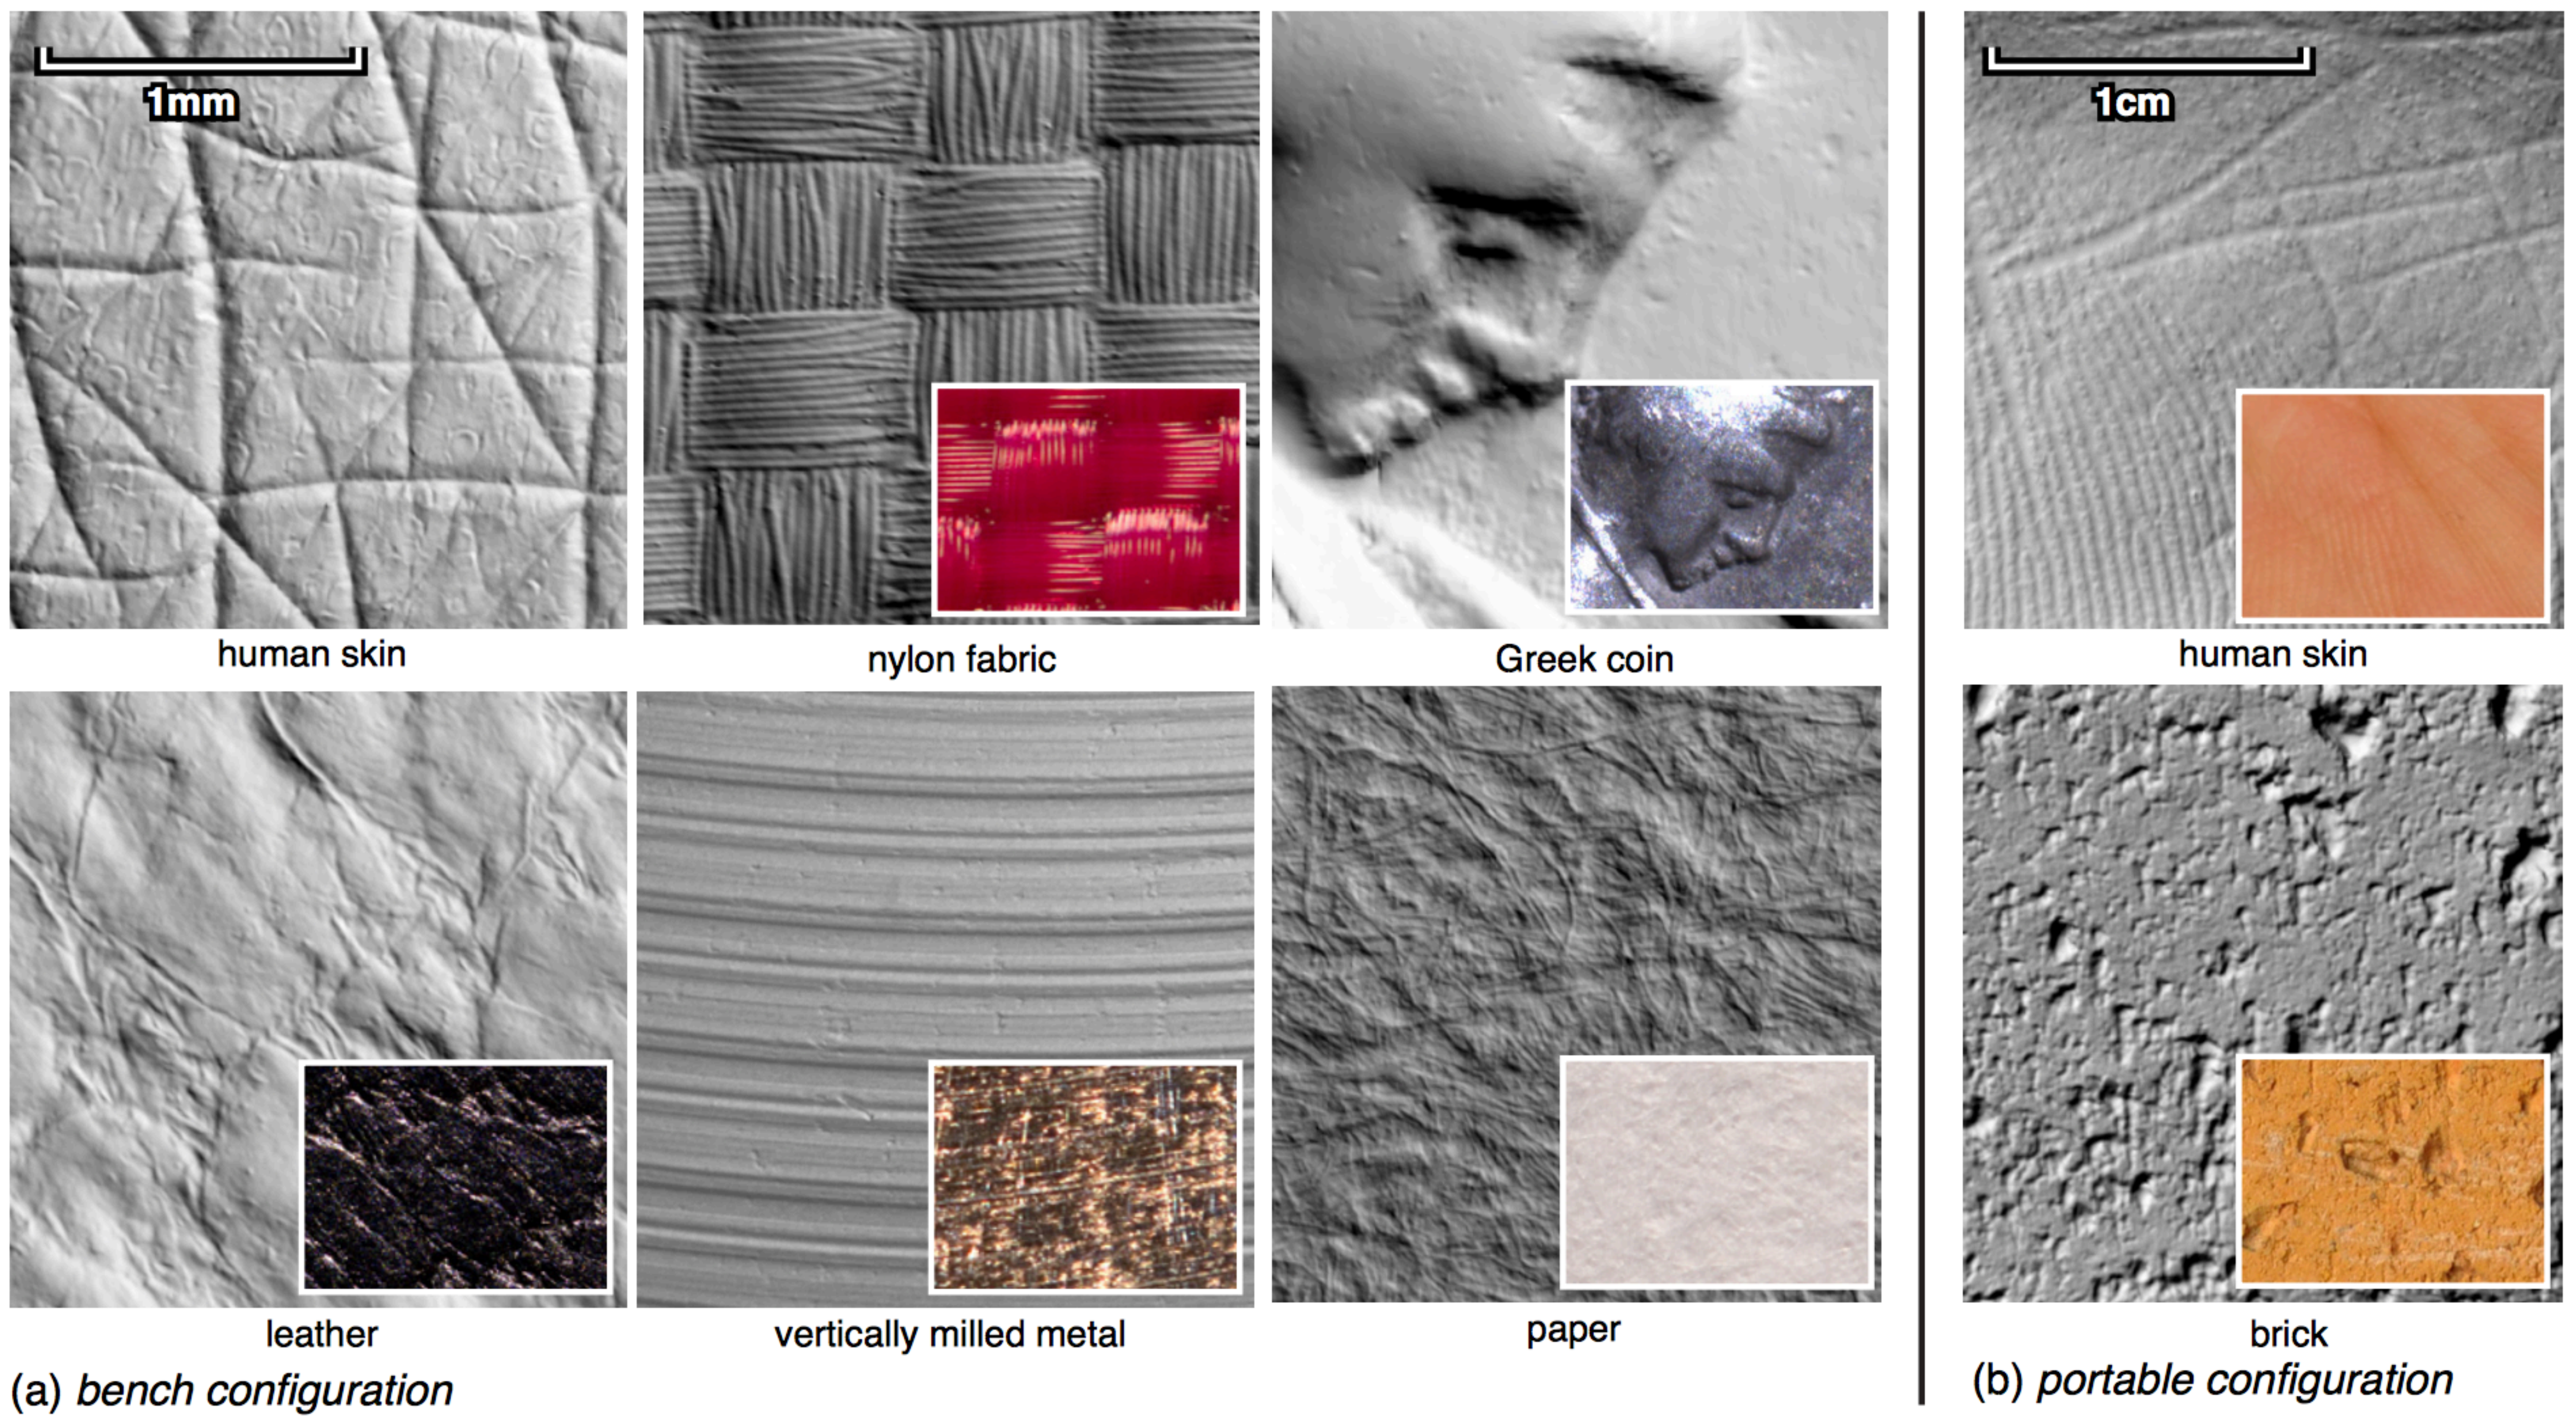

Figure 9: *Example geometry measured with the bench and portable configurations. Outer image: rendering under direct lighting. Inset: macro photograph of original sample. Scale shown in upper left. Color images are shown for context and are to similar, but not exact scale.*

What about paint?

$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

k_d is **reflectance** or **albedo**

Intrinsic image decomposition

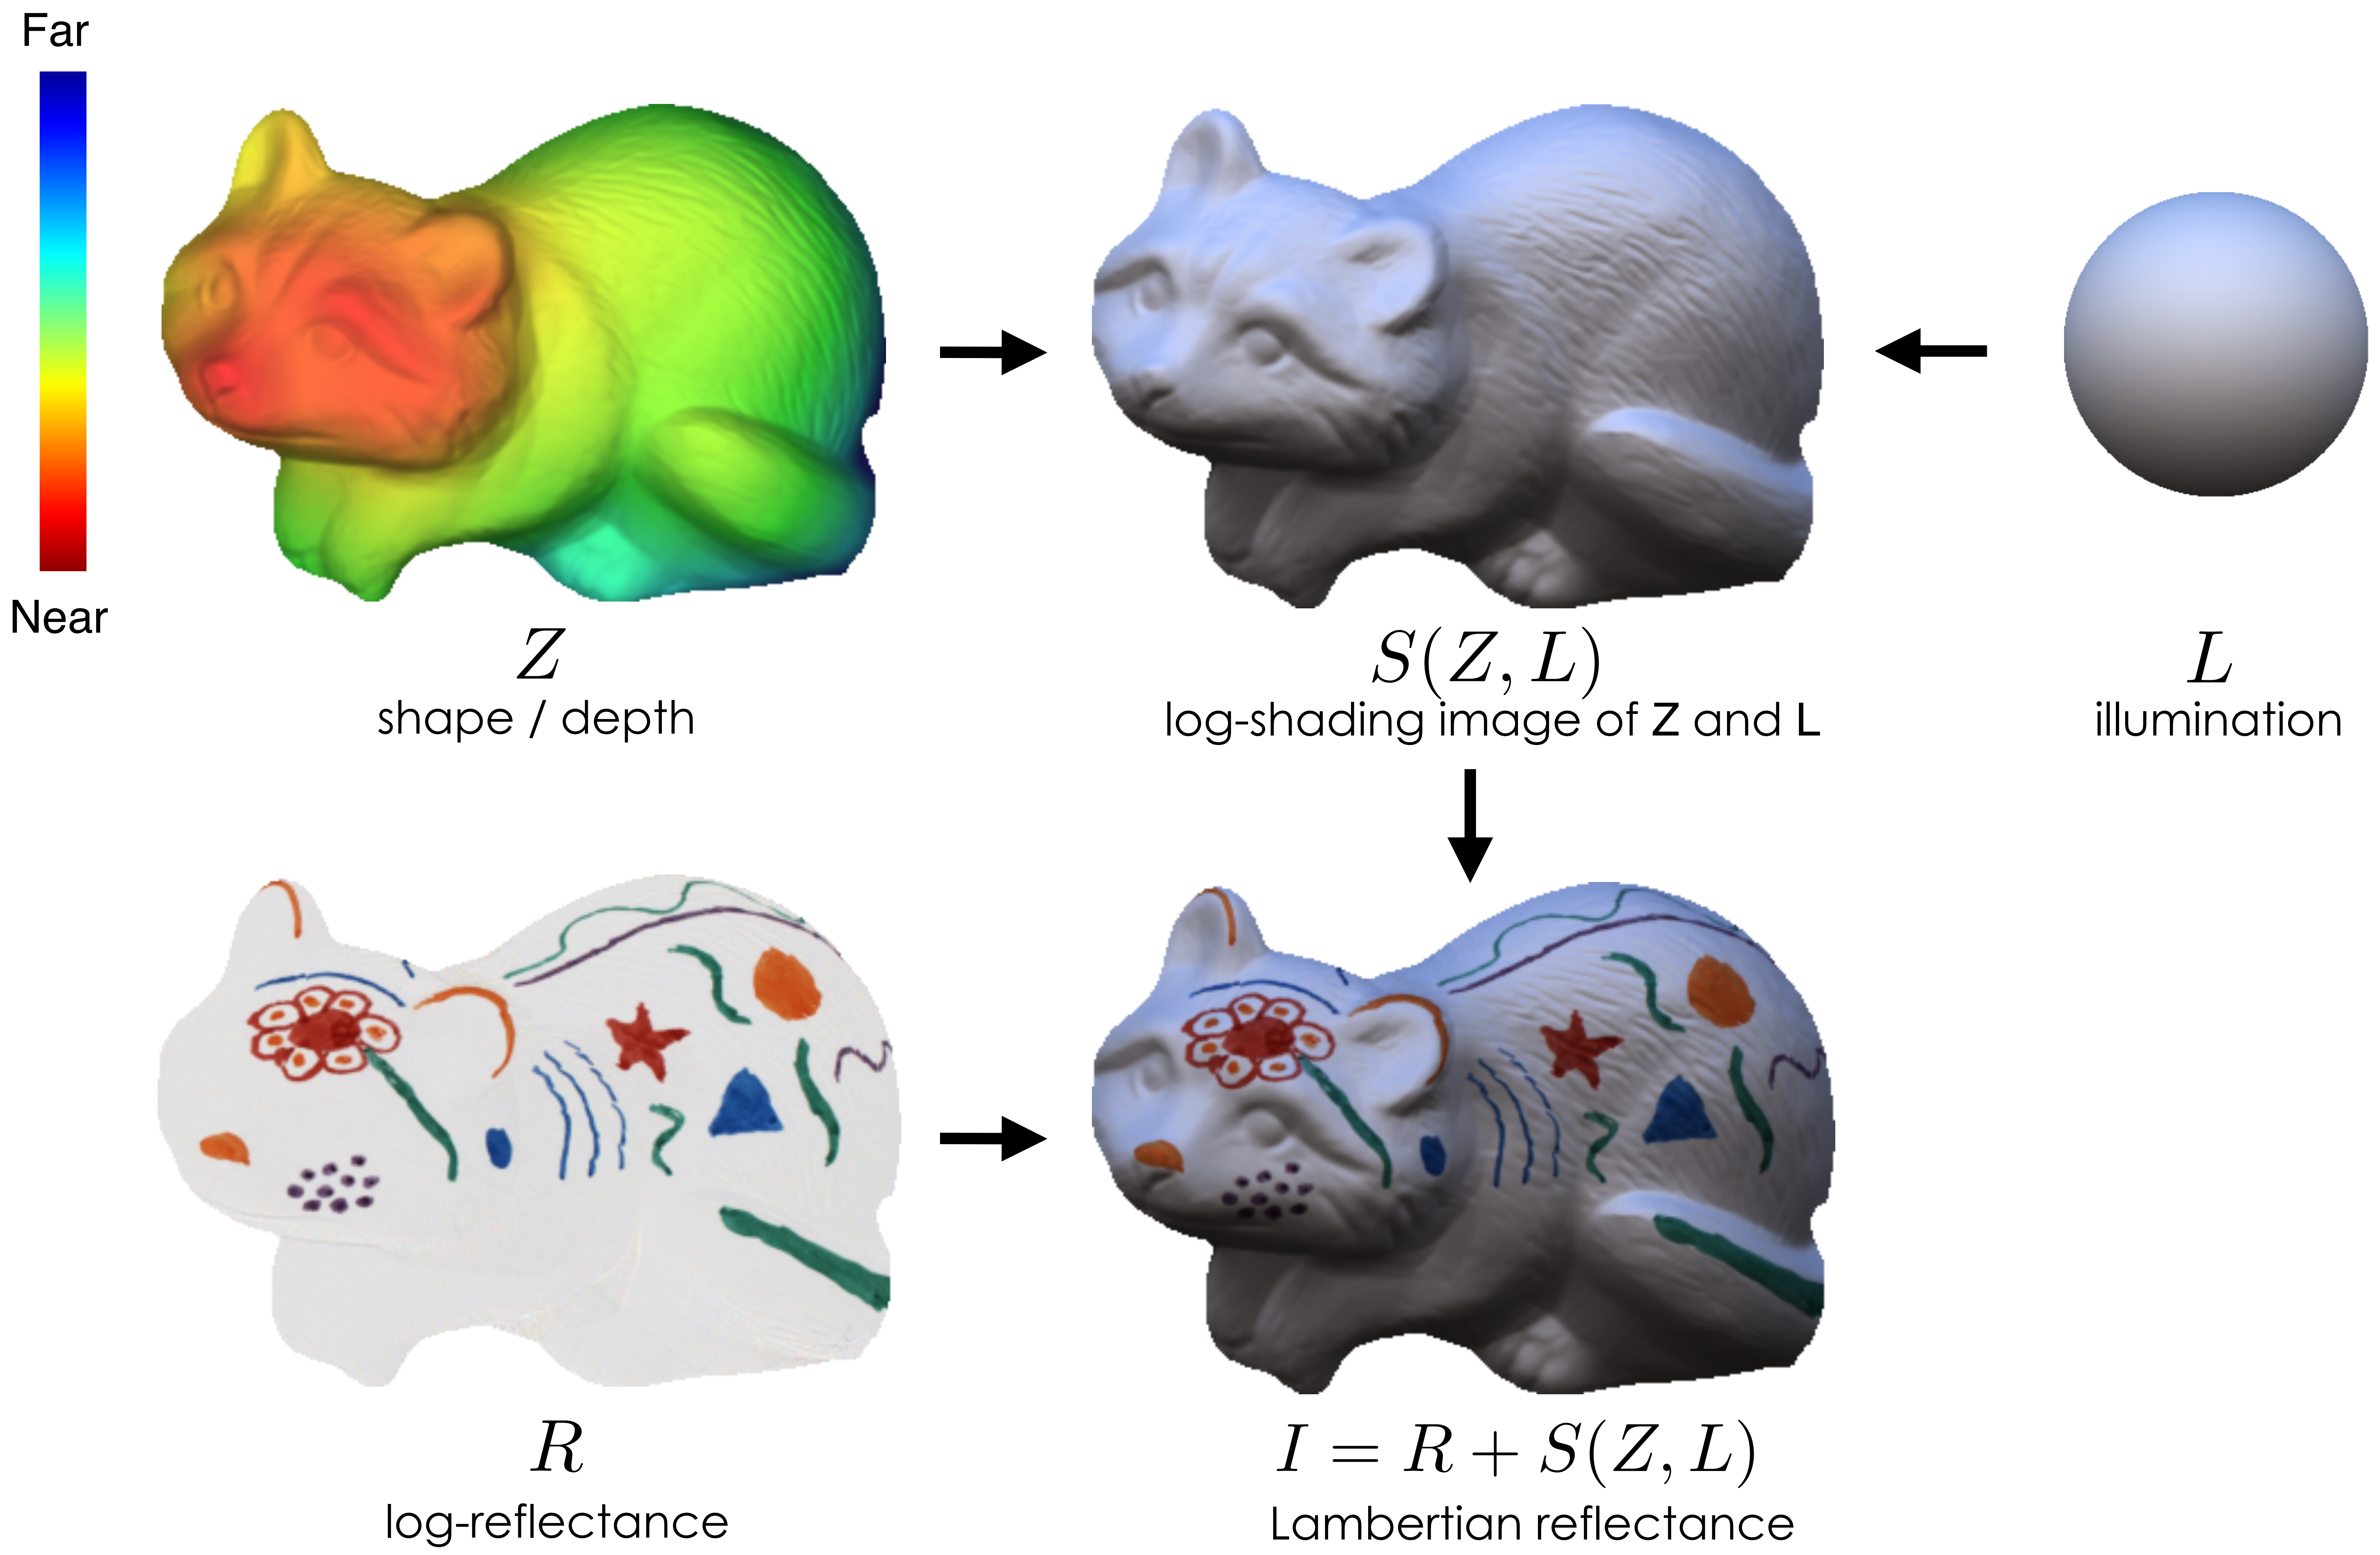

Source: J. Barron

Intrinsic image decomposition

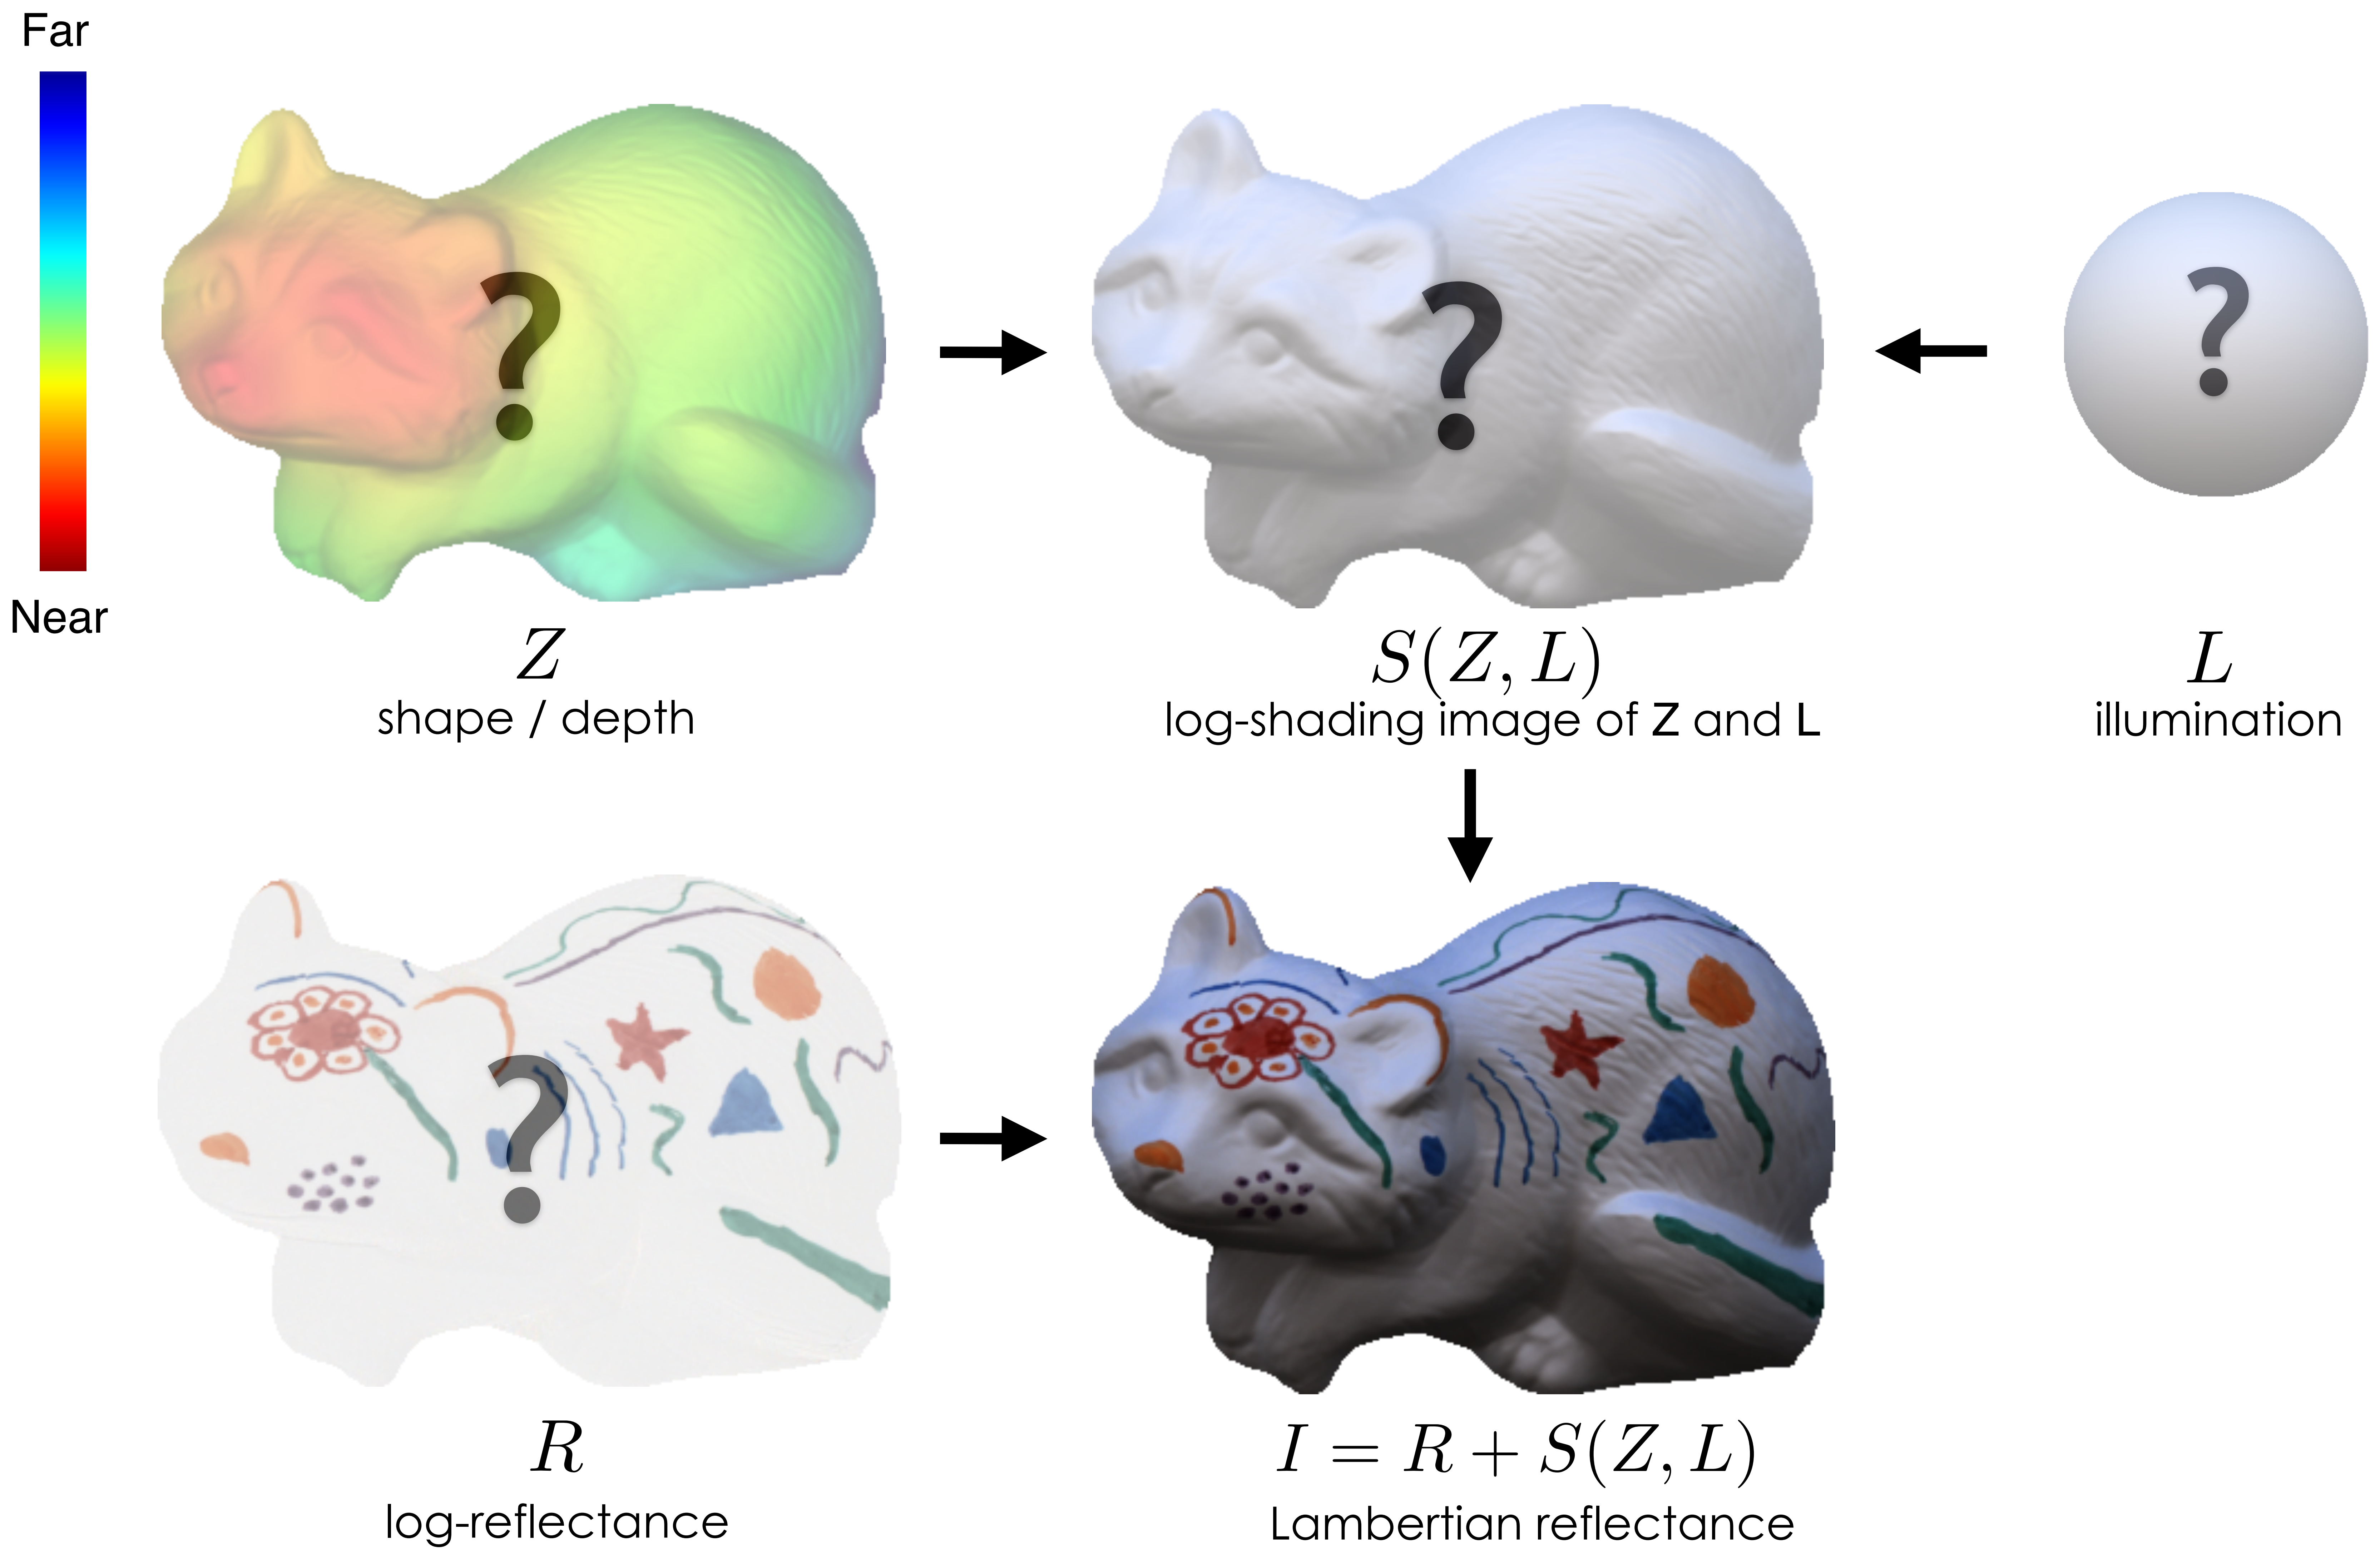

Source: J. Barron

Intrinsic image decomposition

Reflectance

Shading

CNN-based reflectance estimation

Input

Reflectance

Shading

[Bell et al., “Intrinsic images in the wild”, 2014]

Applications of intrinsic image decomposition

[Barron and Malik “SIRFS”, 2012]

Application: relighting

[Barron and Malik “Scene-SIRFS”, 2013]

Application: relighting

[Barron and Malik “Scene-SIRFS”, 2013]

Next week: perceptual grouping