AI Art at Christie's Sells for $432,500

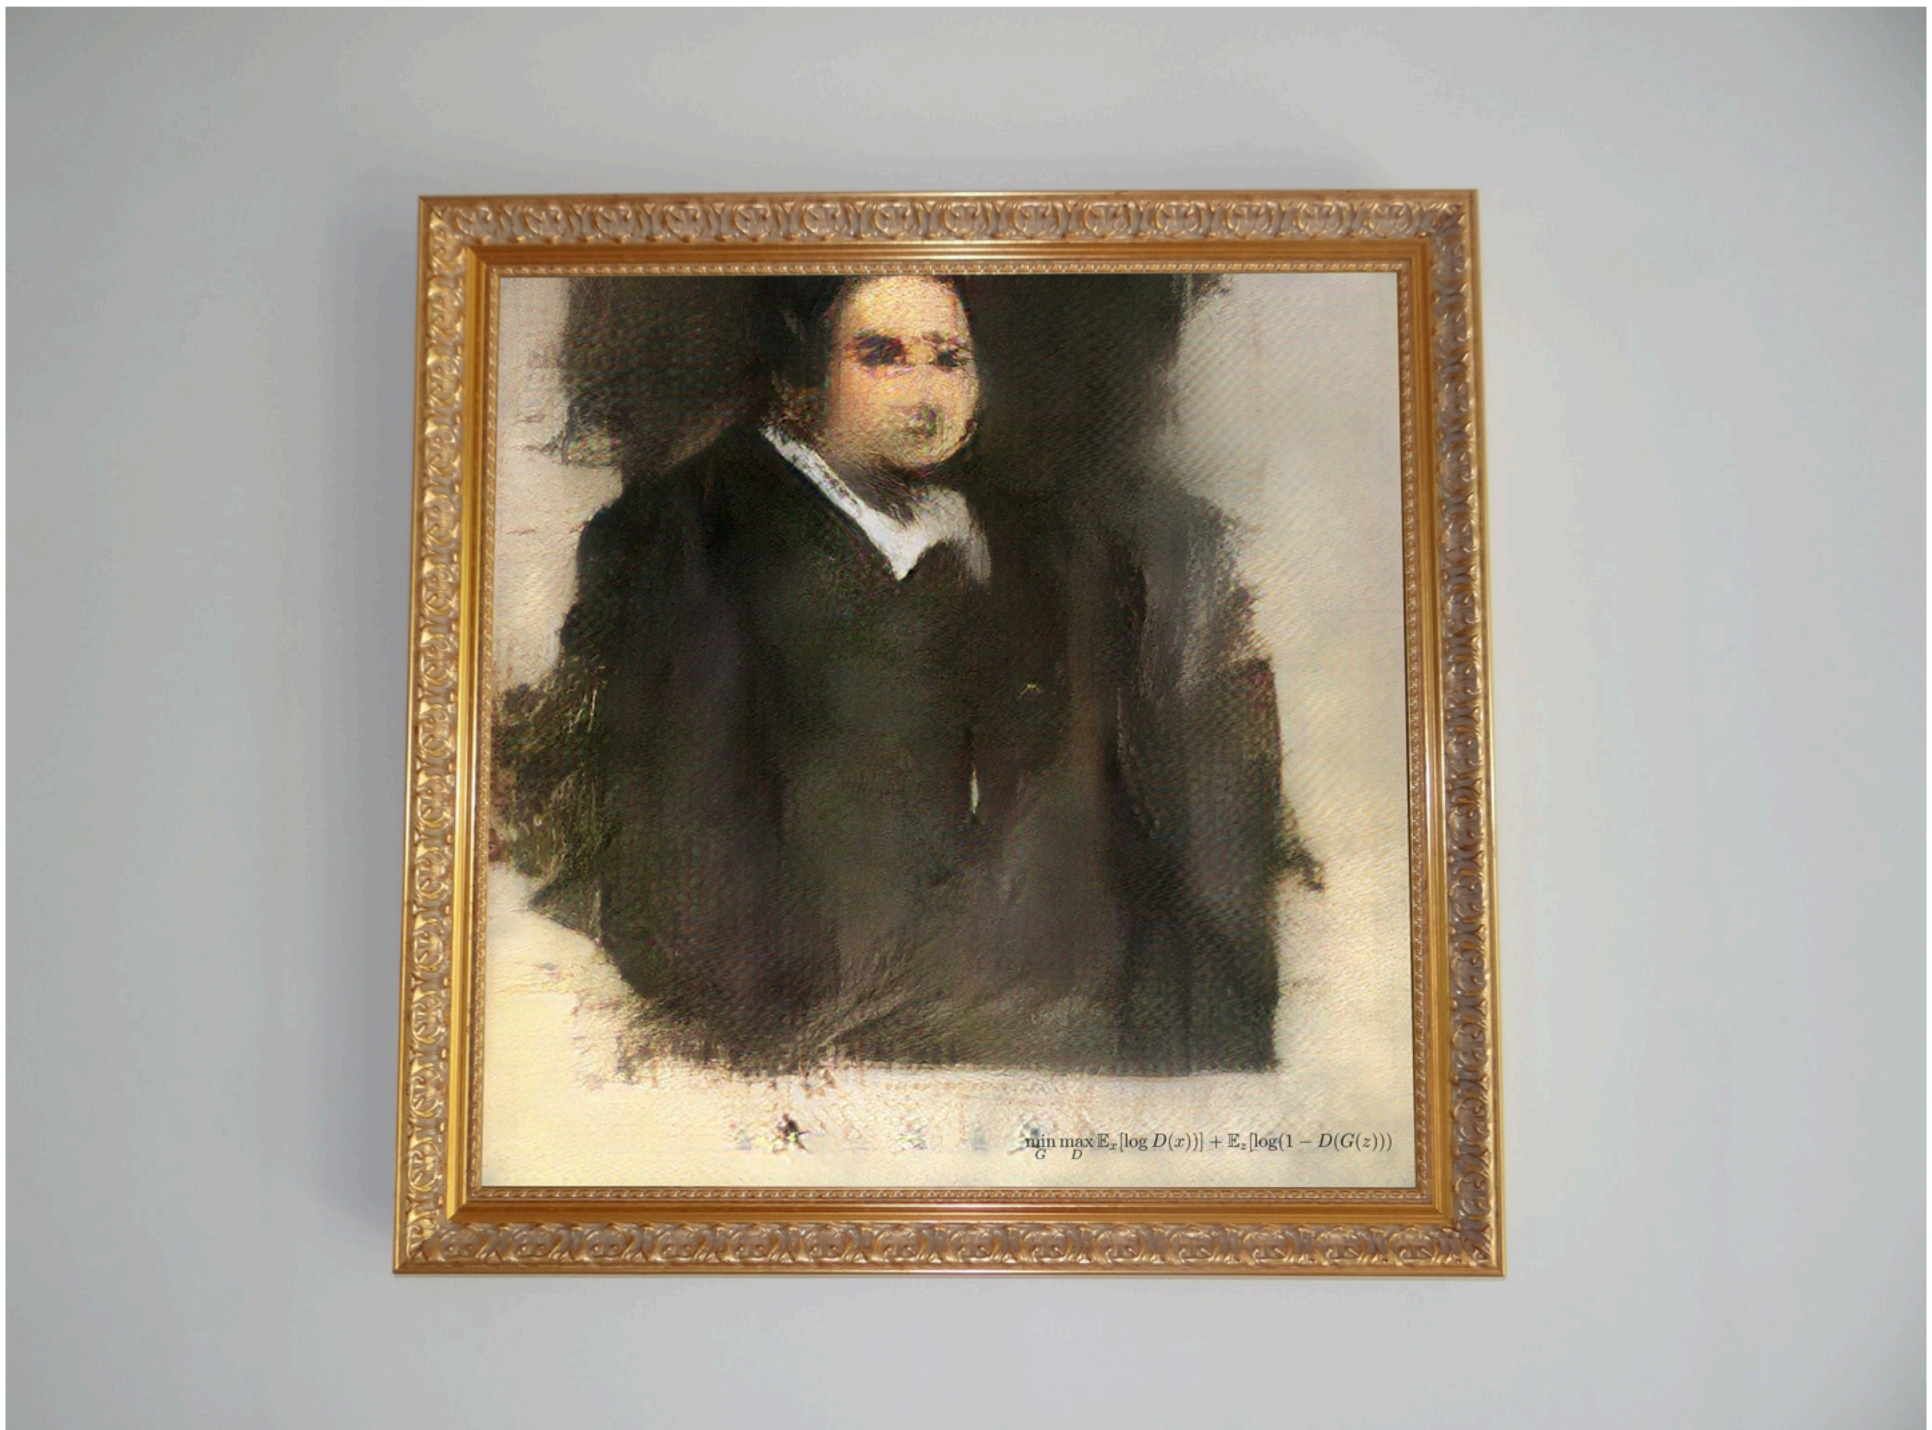

Source: NY Times

Photometric Stereo

Computer Vision
Fall 2018
Columbia University

Last Time: Two-View Stereo

Key Idea: use feature motion to understand shape

Today: Photometric Stereo

Key Idea: use pixel brightness to understand shape

Today: Photometric Stereo

Key Idea: use pixel brightness to understand shape

Foreshortening

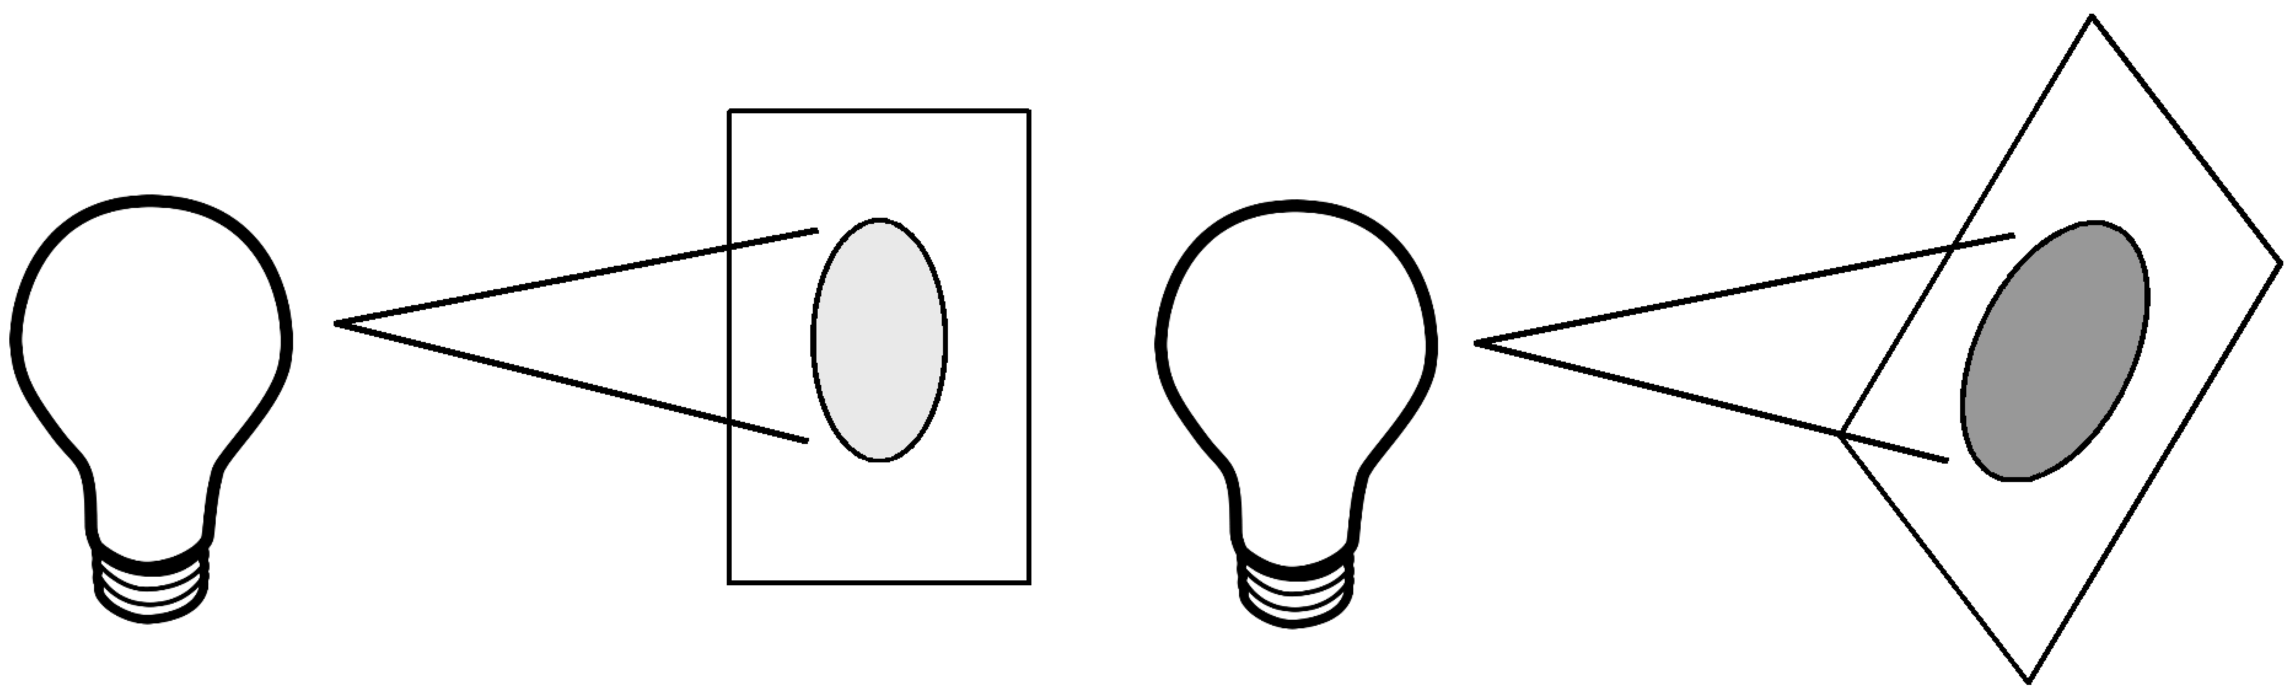

As the surface tilts away from the light source the same light energy is spread over a larger area, making the surface darker

Photometric Stereo

What results can you get?

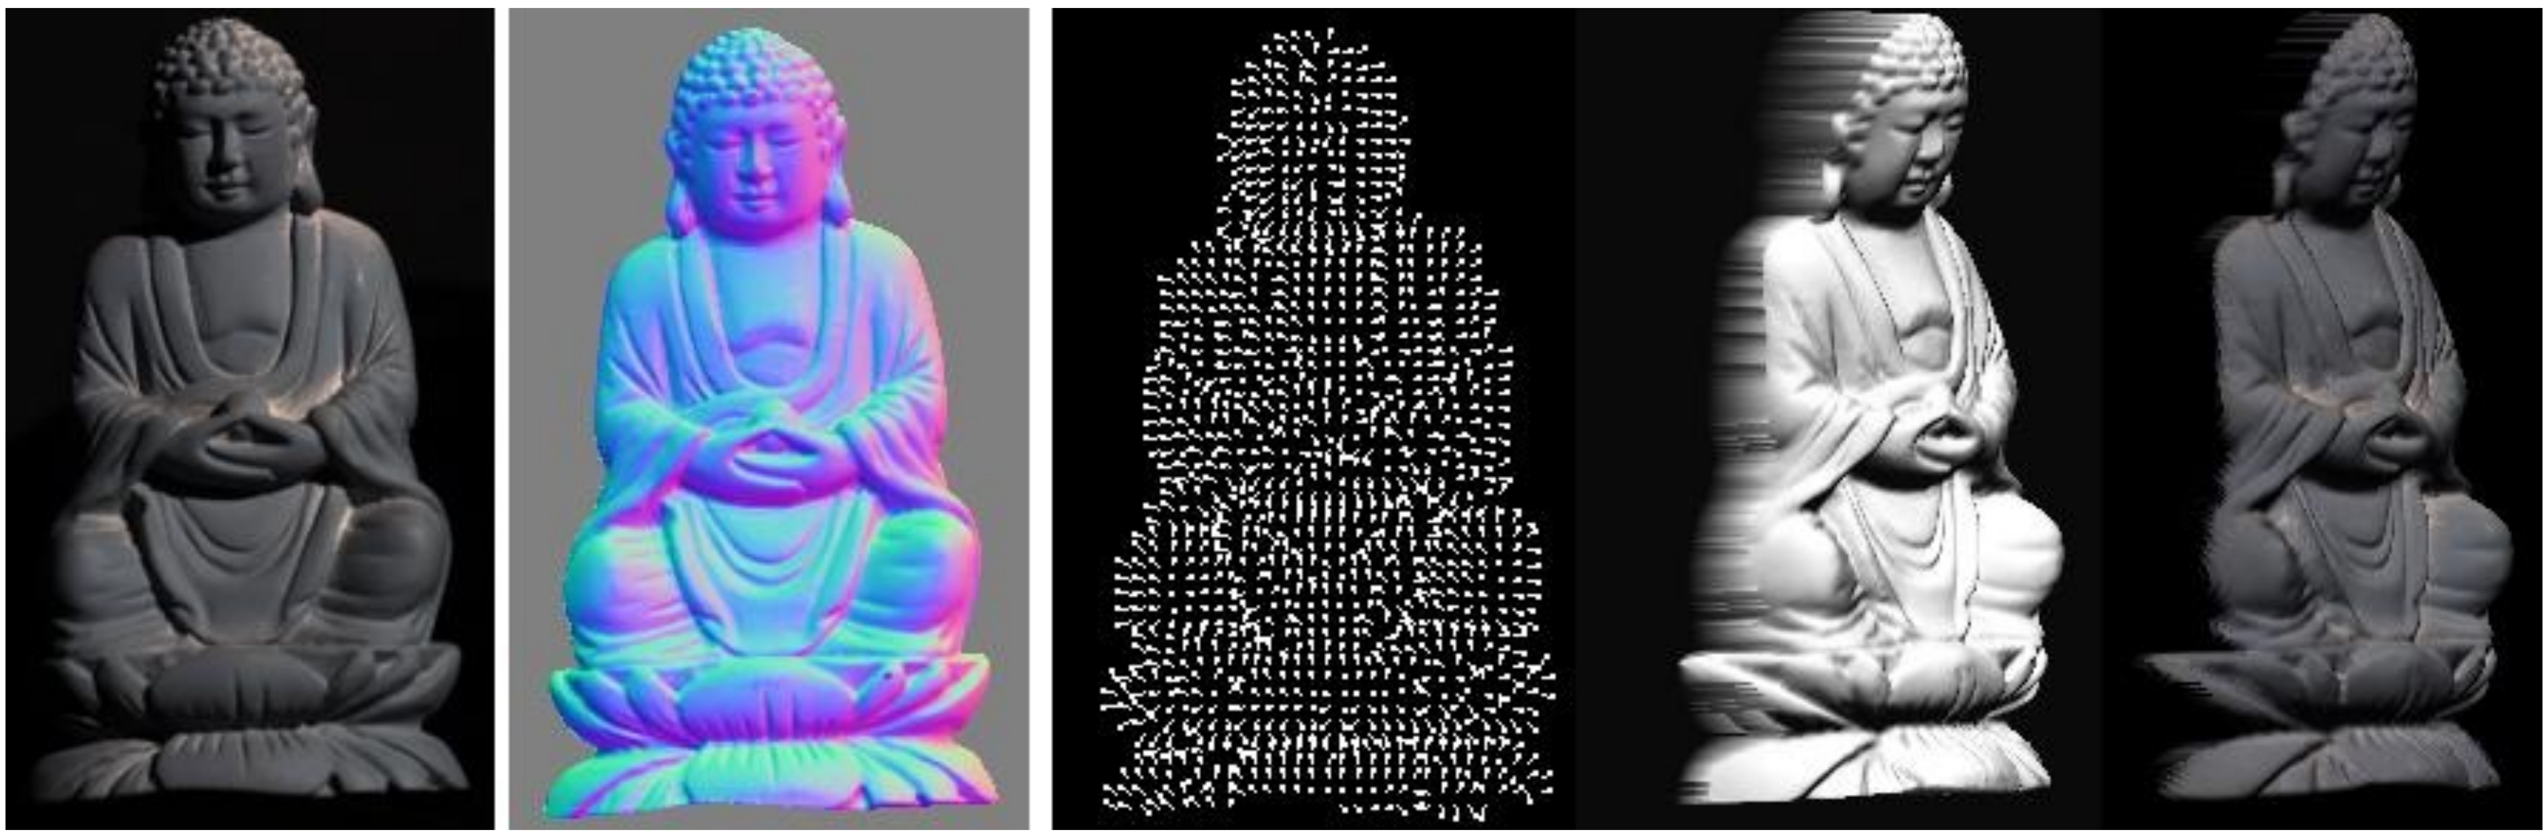

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

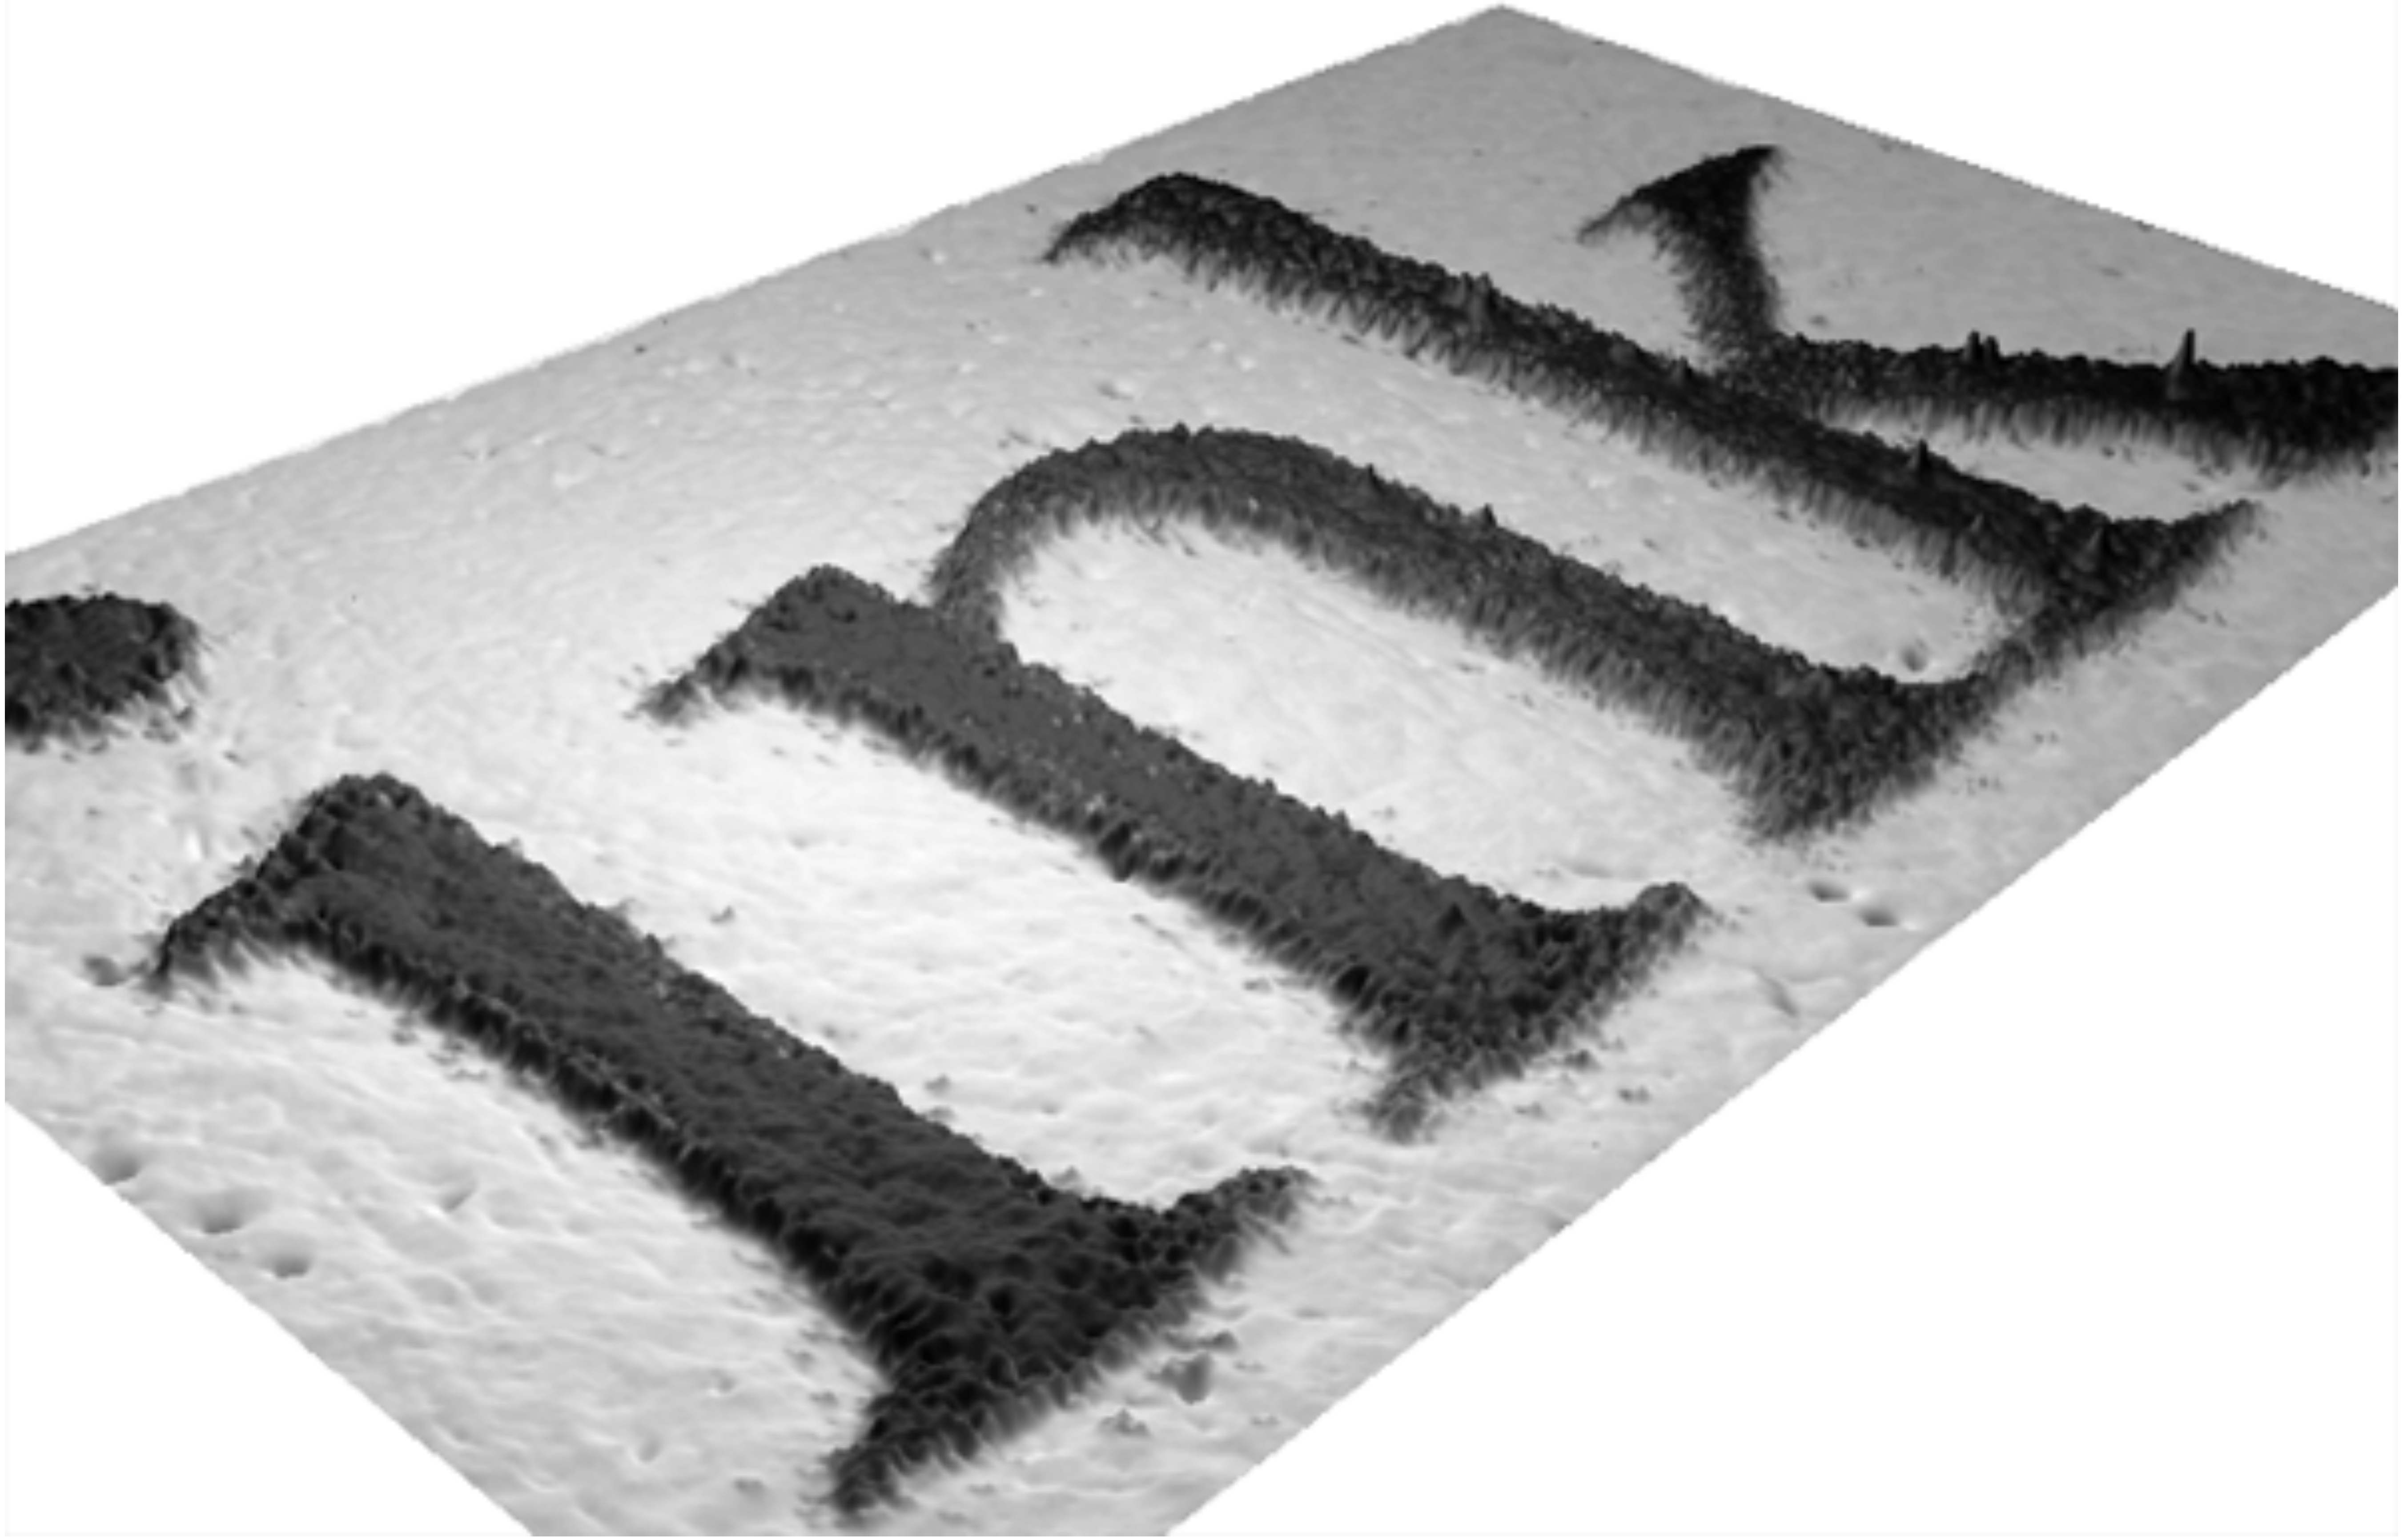
ink

Modeling Image Formation

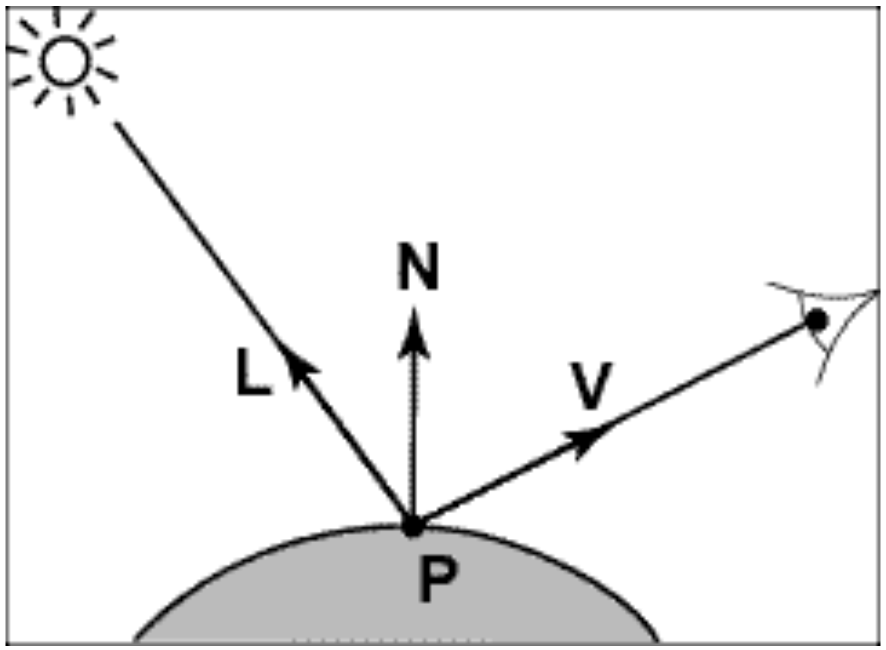

Now we need to reason about:

- How light interacts with the scene
- How a pixel value is related to light energy in the world

Track a “ray” of light all the way from light source to the sensor

Light rays interacting with a surface

- Light of radiance L_i comes from light source at an incoming direction θ_i
- It sends out a ray of radiance L_r in the outgoing direction θ_r
- How does L_r relate to L_i ?

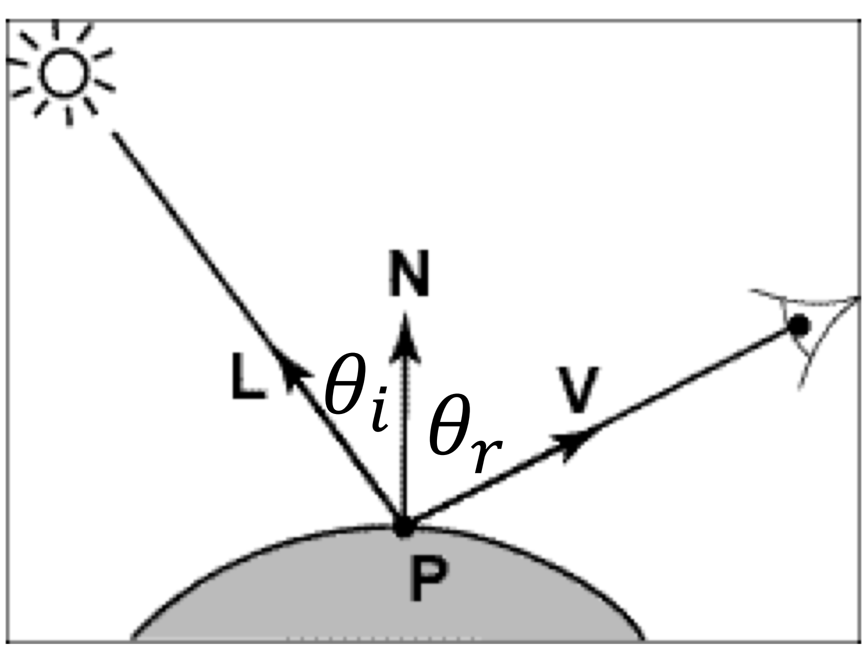

- **N** is surface normal
- **L** is direction of light, making θ_i with normal
- **V** is viewing direction, making θ_r with normal

Light rays interacting with a surface

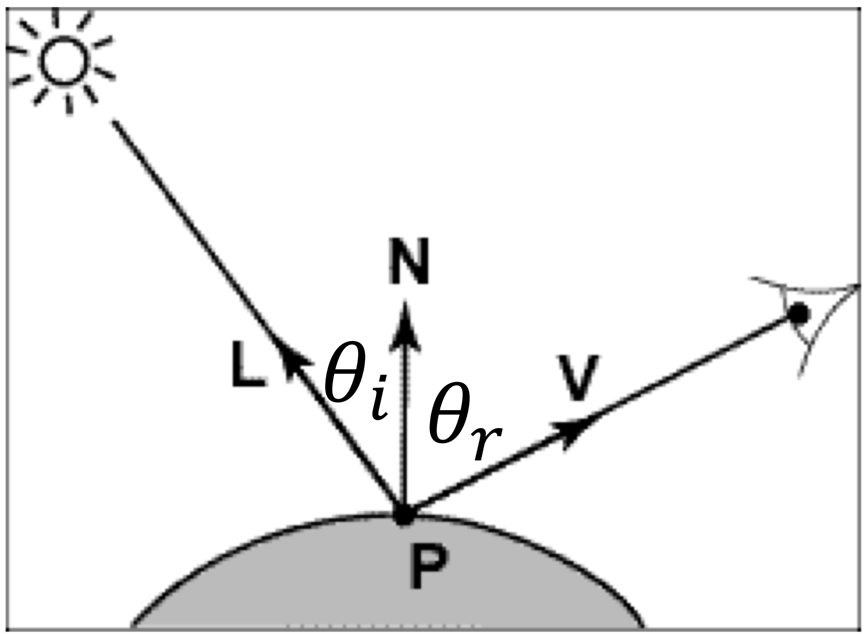

- **N** is surface normal
- **L** is direction of light, making θ_i with normal
- **V** is viewing direction, making θ_r with normal

Output radiance along **V**

$$\boxed{L_r} = \boxed{\rho(\theta_i, \theta_r)}\boxed{L_i \cos\theta_i}$$

Incoming irradiance along **L**

Bi-directional reflectance function (BRDF)

Materials - Three Forms

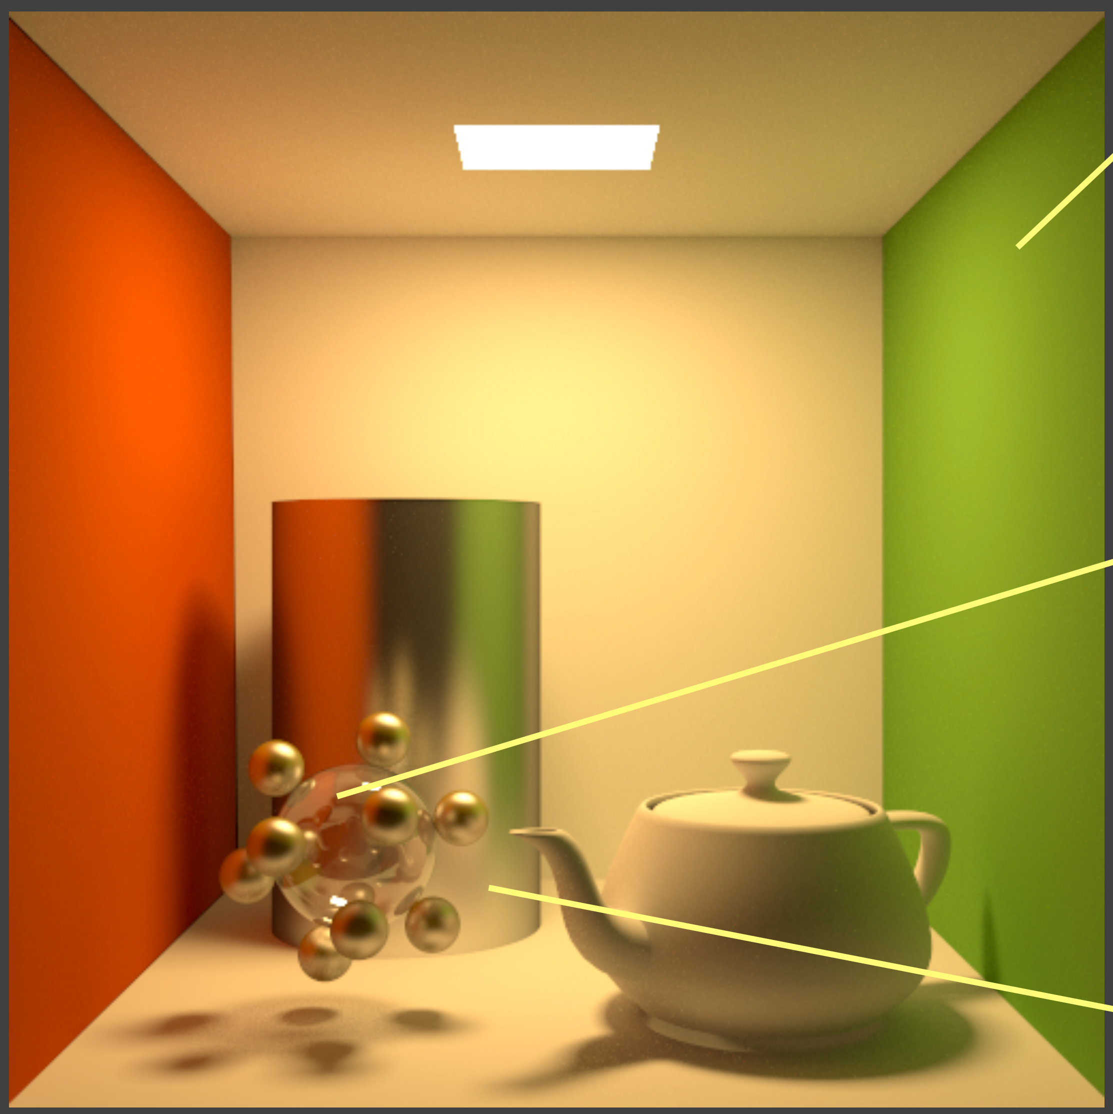

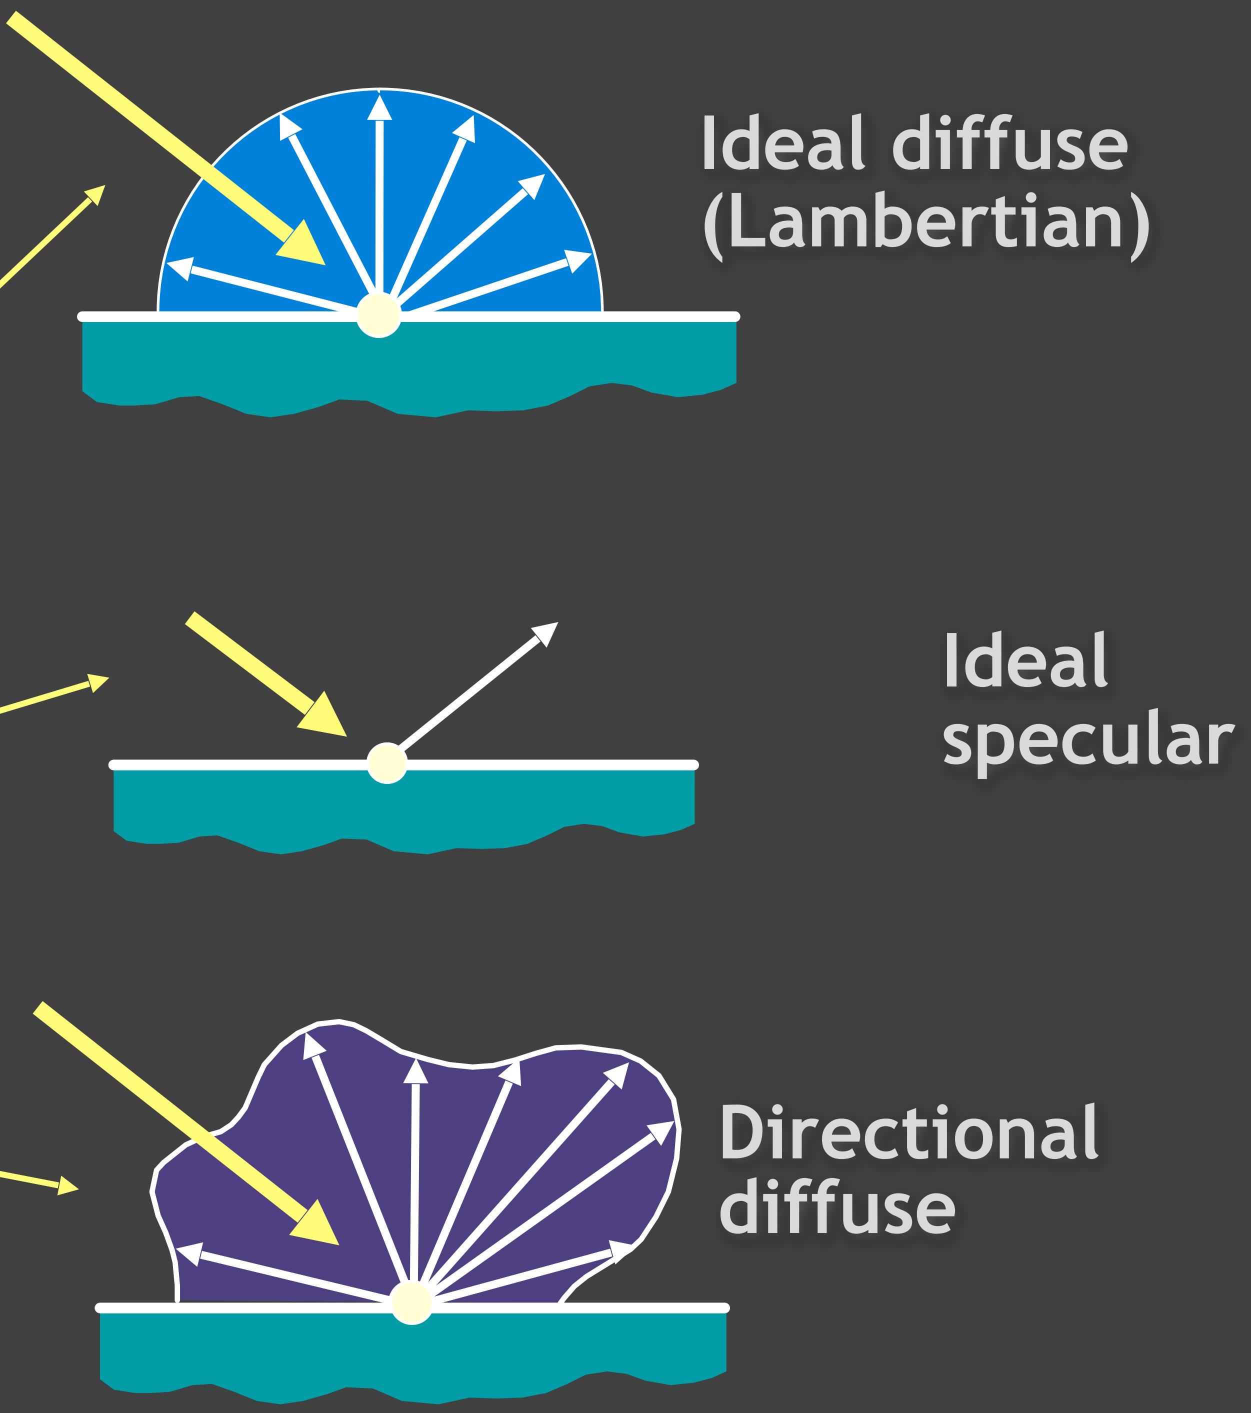

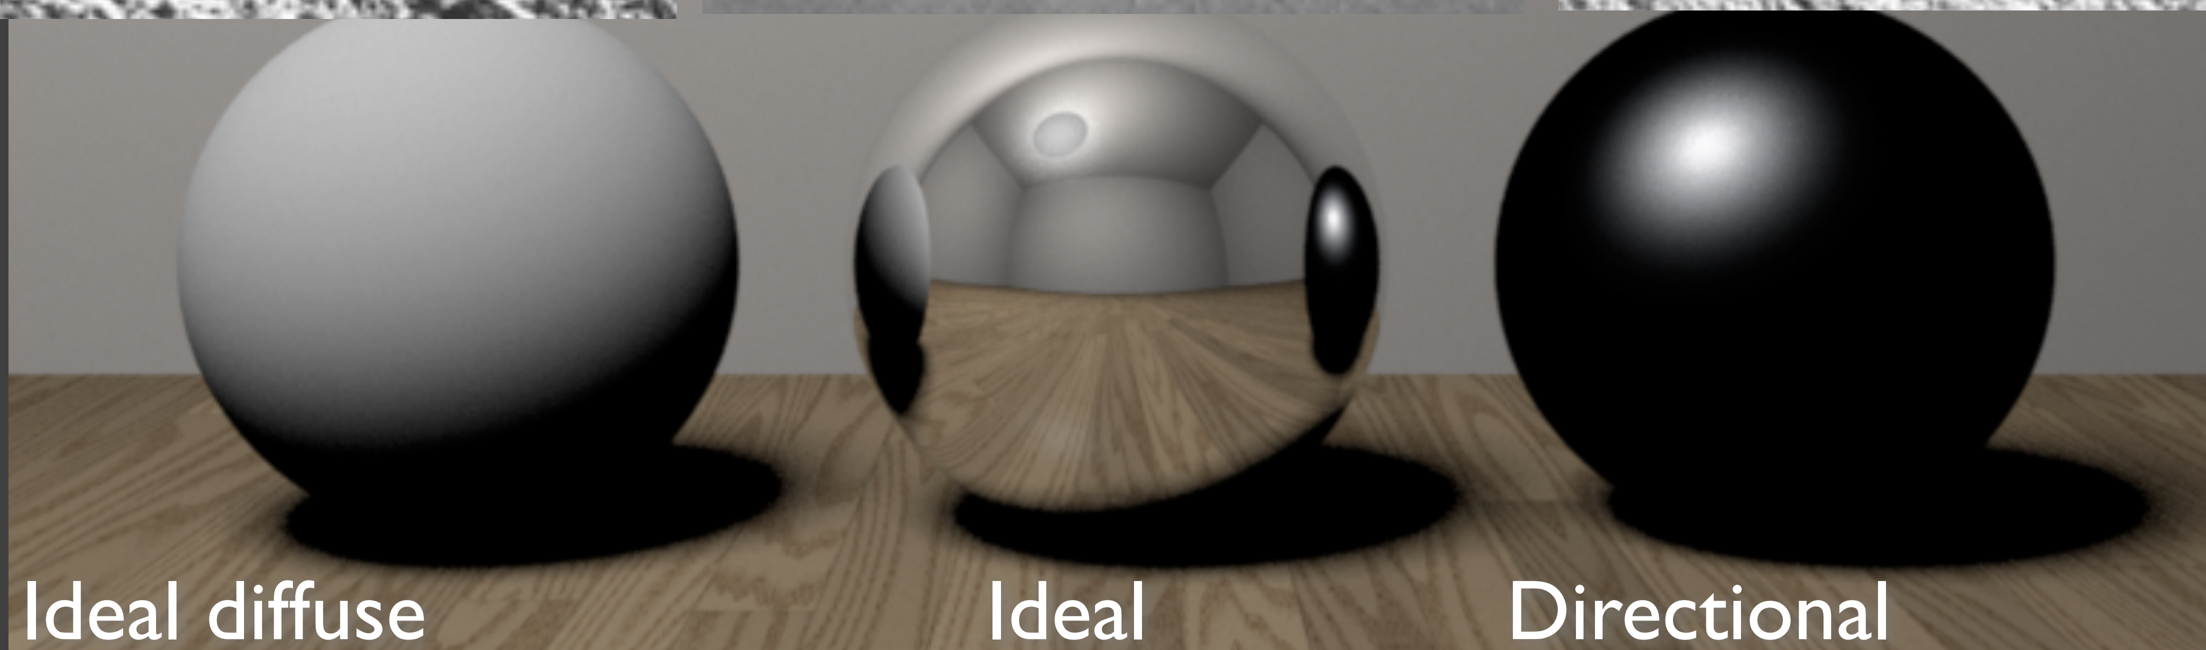

Ideal diffuse (Lambertian)

Ideal specular

Directional diffuse

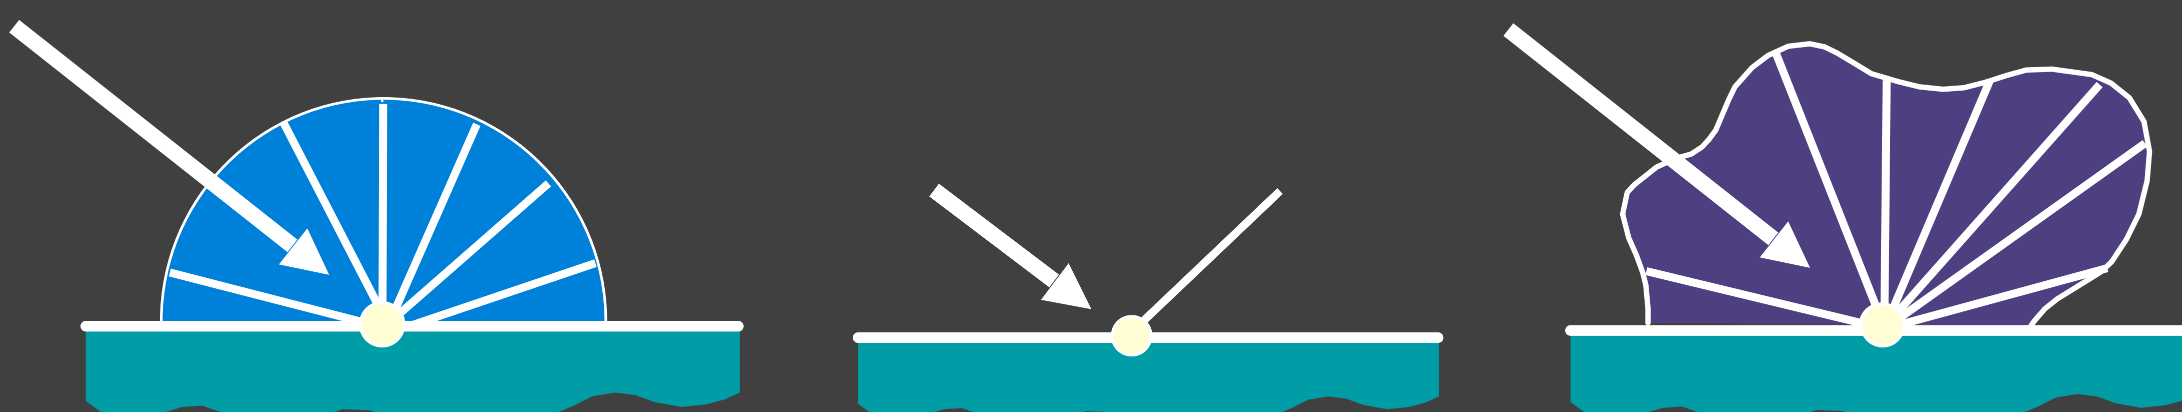

Ideal Diffuse Reflection

- Characteristic of multiple scattering materials
- An idealization but reasonable for matte surfaces
- Basis of most radiosity methods

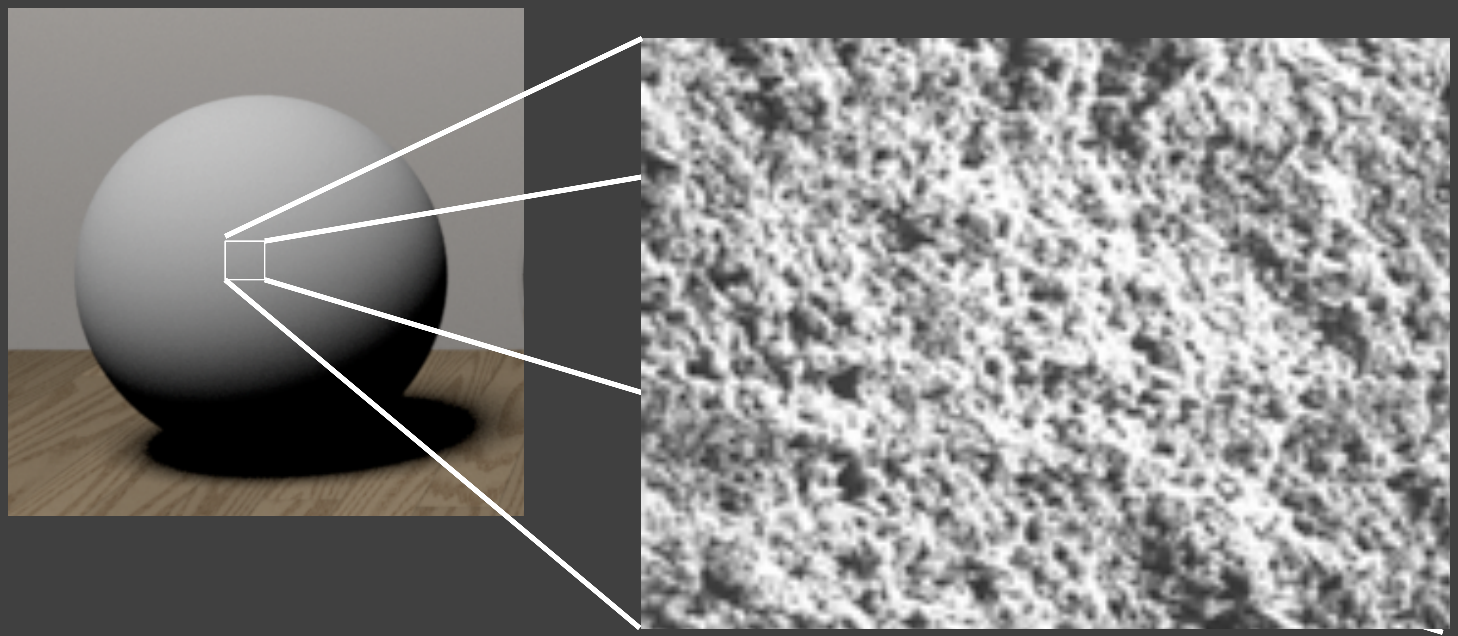

Ideal Specular Reflection

- Calculated from Fresnel’s equations
- Exact for polished surfaces
- Basis of early ray-tracing methods

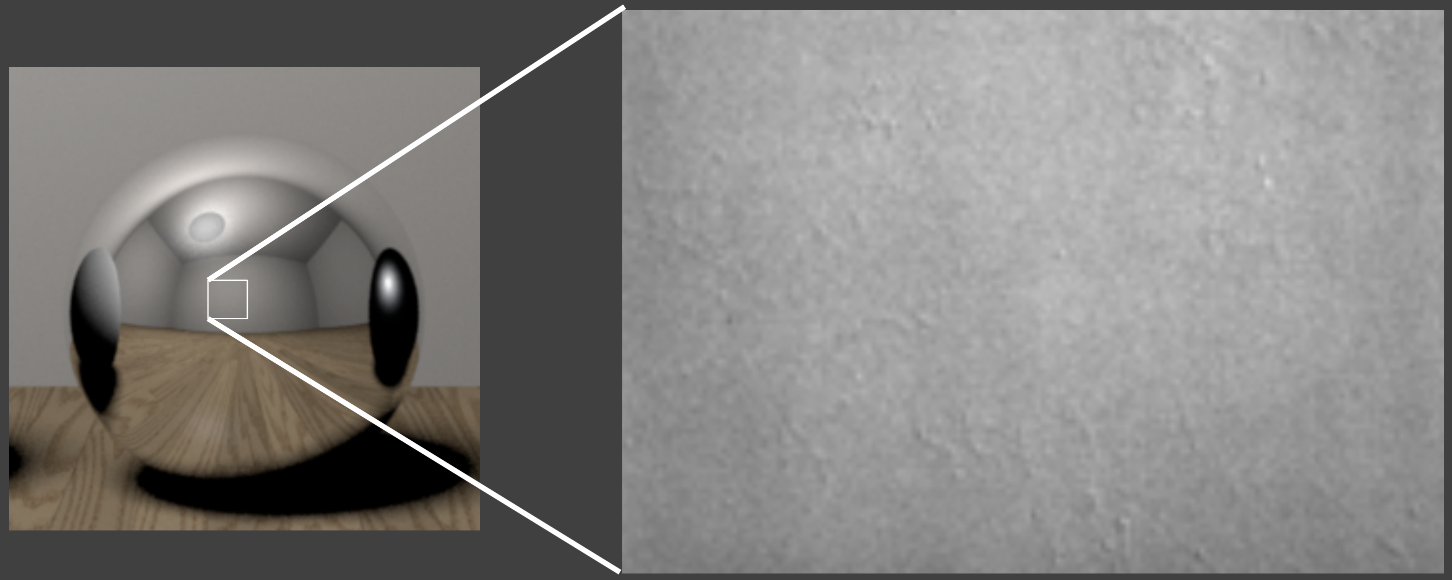

Light rays interacting with a surface

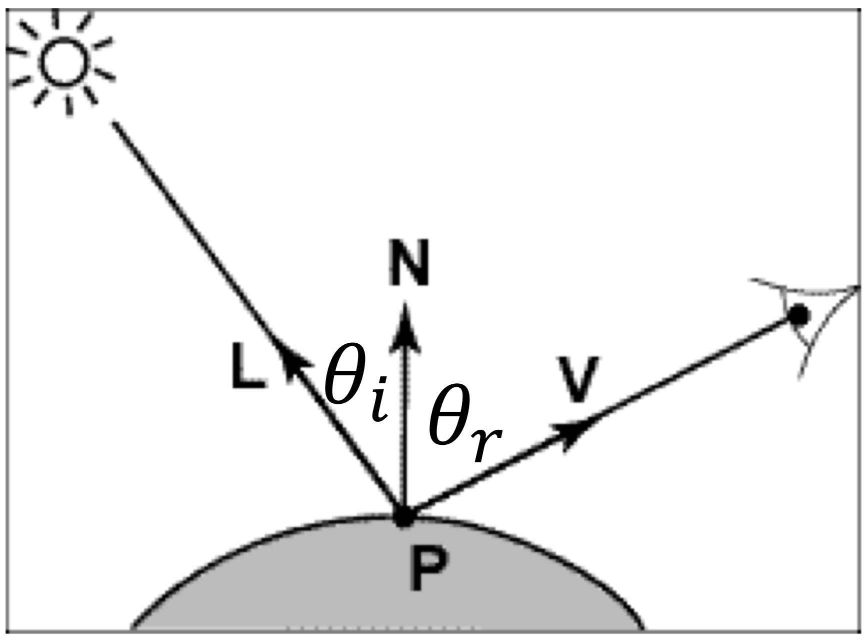

$$L_r = \rho(\theta_i, \theta_r) L_i \cos\theta_i$$

- Special case 1: Perfect mirror
  - $\rho(\theta_i, \theta_r)$ = 0 unless $\theta_i = \theta_r$
- Special case 2: Matte surface
  - $\rho(\theta_i, \theta_r)$ = ρ_0 (constant)

For now, ignore specular reflection

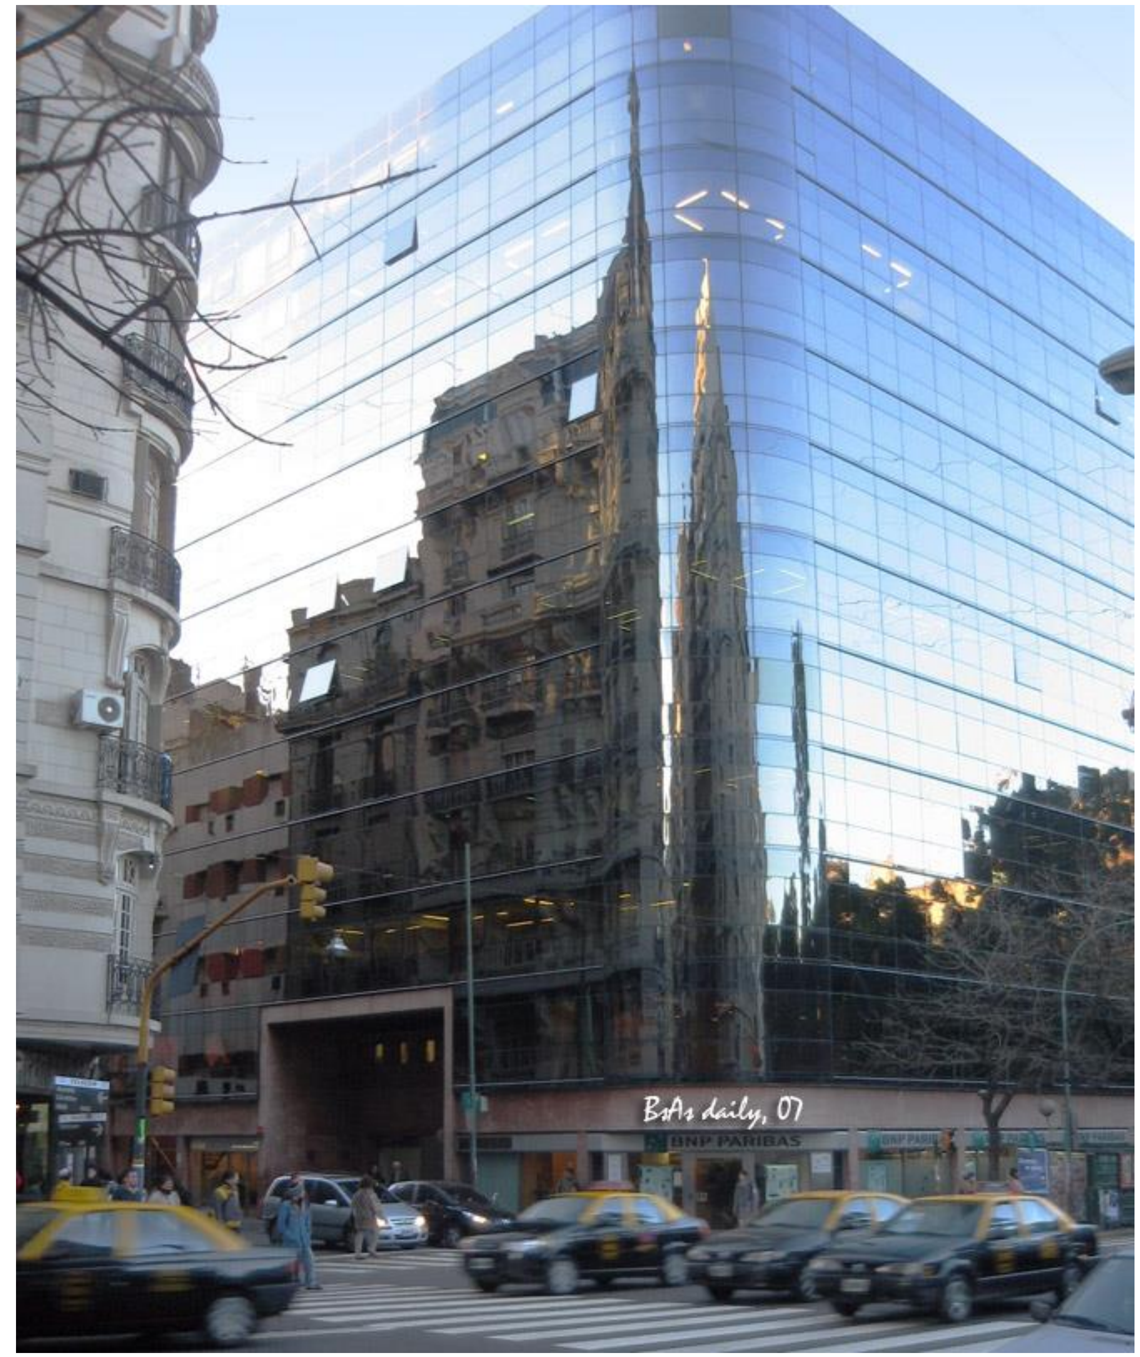

Slides from Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra

And Refraction…

Slides from Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra

And Interreflections…

Slides from Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra

And Subsurface Scattering…

Slides from Photometric Methods for 3D Modeling, Matsushita, Wilburn, Ben-Ezra

Lambertian surfaces

- For a lambertian surface:

$$L_r = \rho L_i \cos\theta_i$$

$$\Rightarrow L_r = \rho L_i \mathbf{L} \cdot \mathbf{N}$$

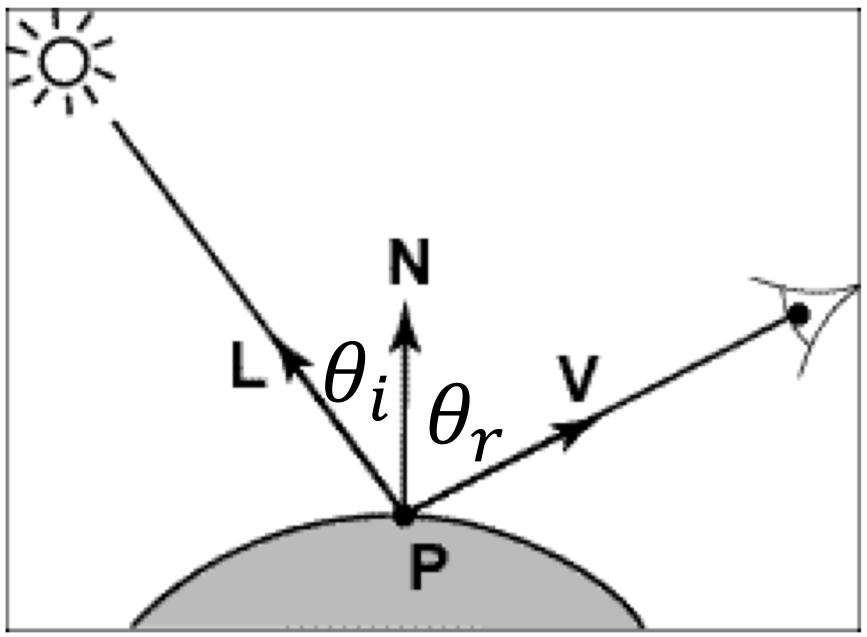

- ρ is called *albedo*
  - Think of this as paint
  - High albedo: white colored surface
  - Low albedo: black surface
  - Varies from point to point

Lambertian surfaces

- Assume the light is directional: all rays from light source are parallel
  - Equivalent to a light source infinitely far away
- All pixels get light from the same direction **L** and of the same intensity L_i

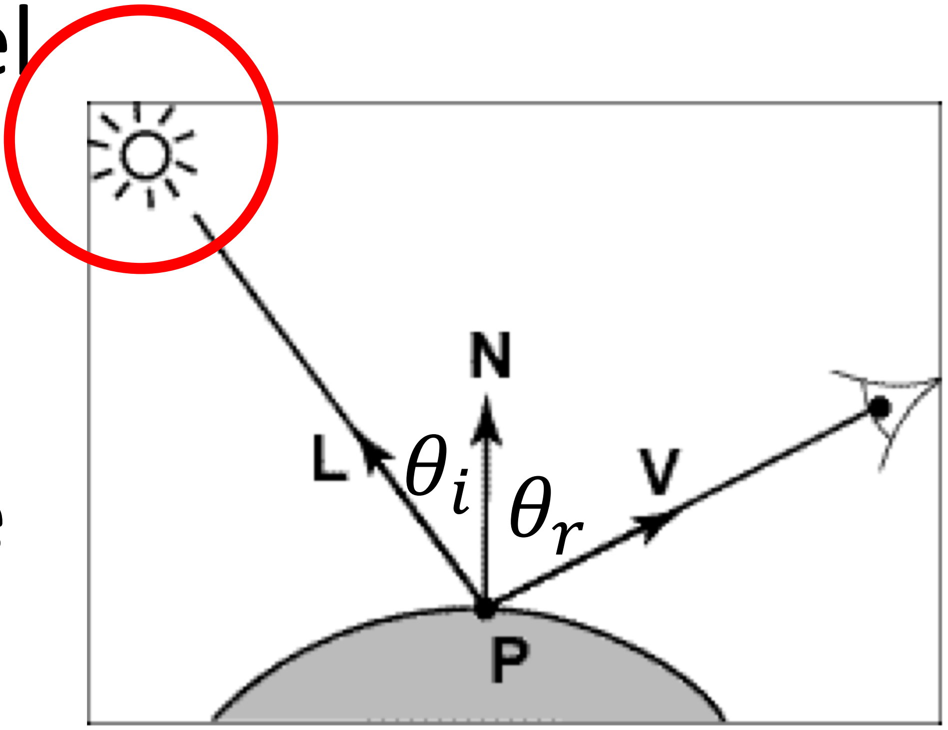

Lambertian surfaces

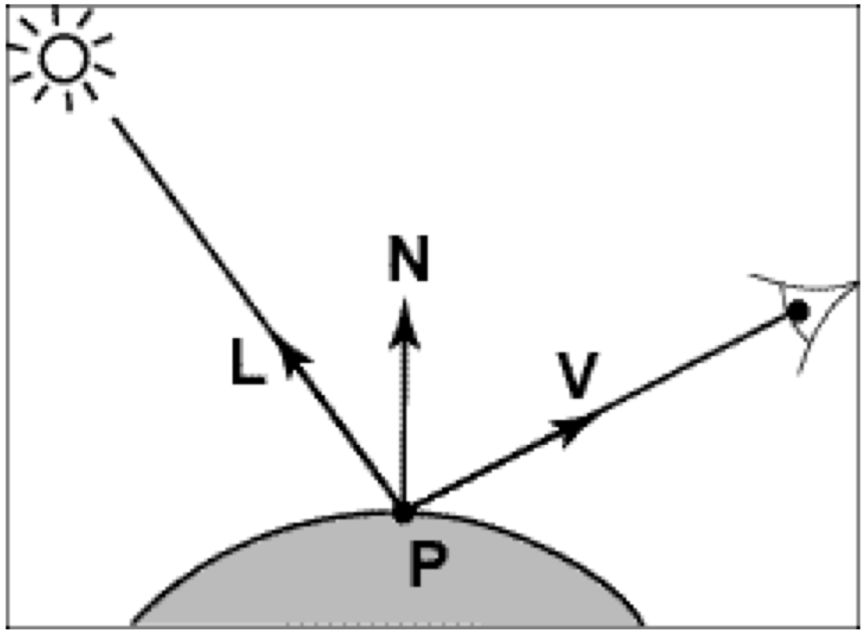

Intrinsic Image Decomposition

$$I(x,y) = \rho(x,y) L_i \mathbf{L} \cdot \mathbf{N}(x,y)$$

Reflectance image

Shading image

Reconstructing Lambertian surfaces

$$I(x, y) = \rho(x, y) L_i \mathbf{L} \cdot \mathbf{N}(x, y)$$

- Equation is a constraint on albedo and normals
- Can we solve for albedo and normals?

Solution 1: Recovery from a single image

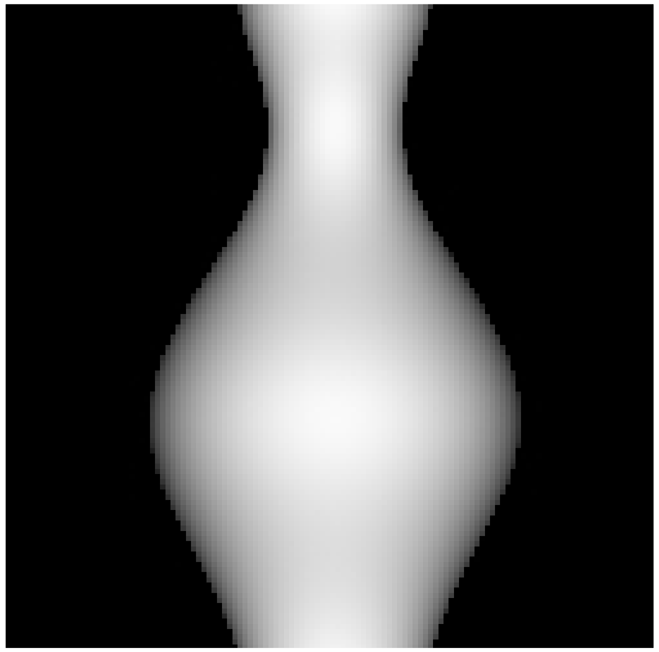  

Solution 1: Recovery from a single image

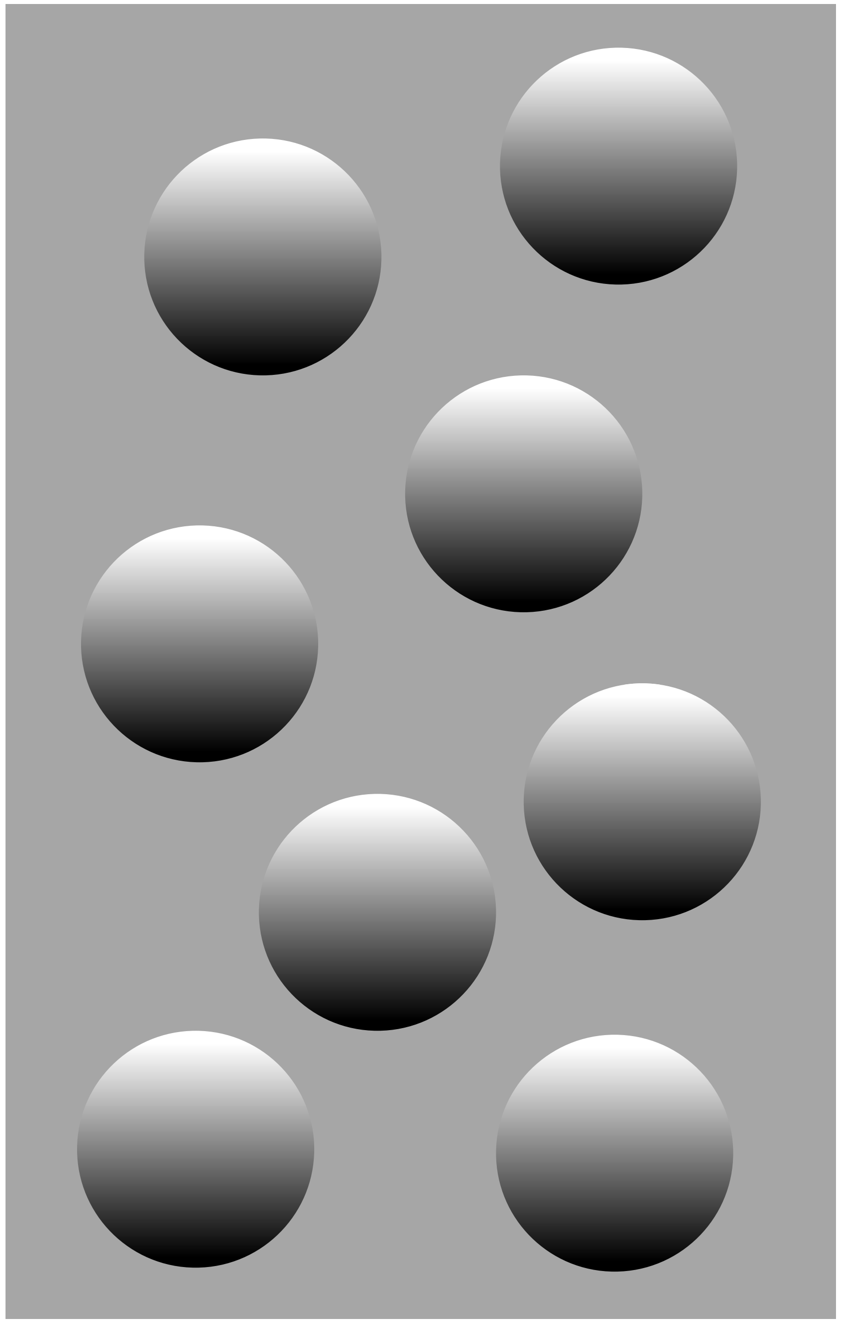

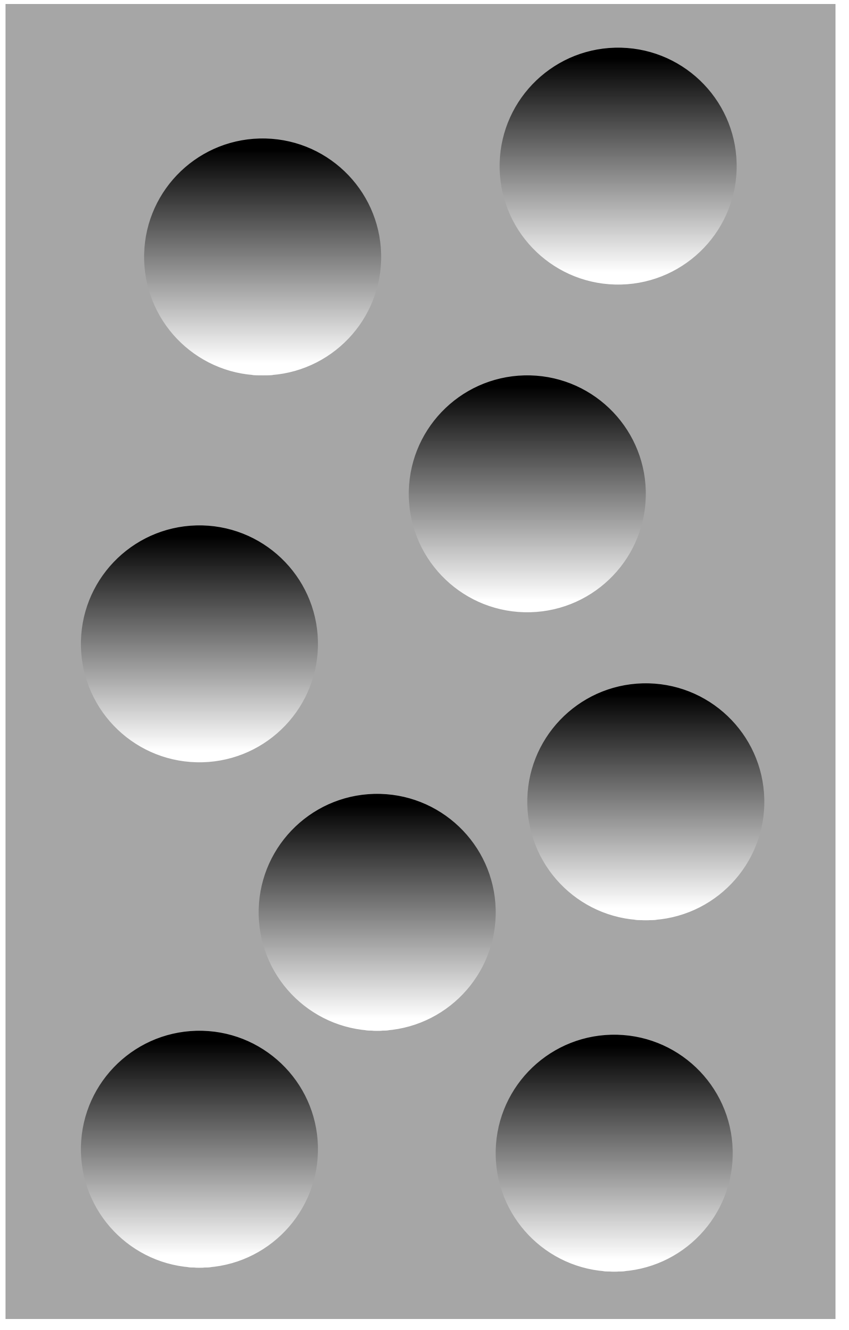

Solution 1: Recovery from a single image

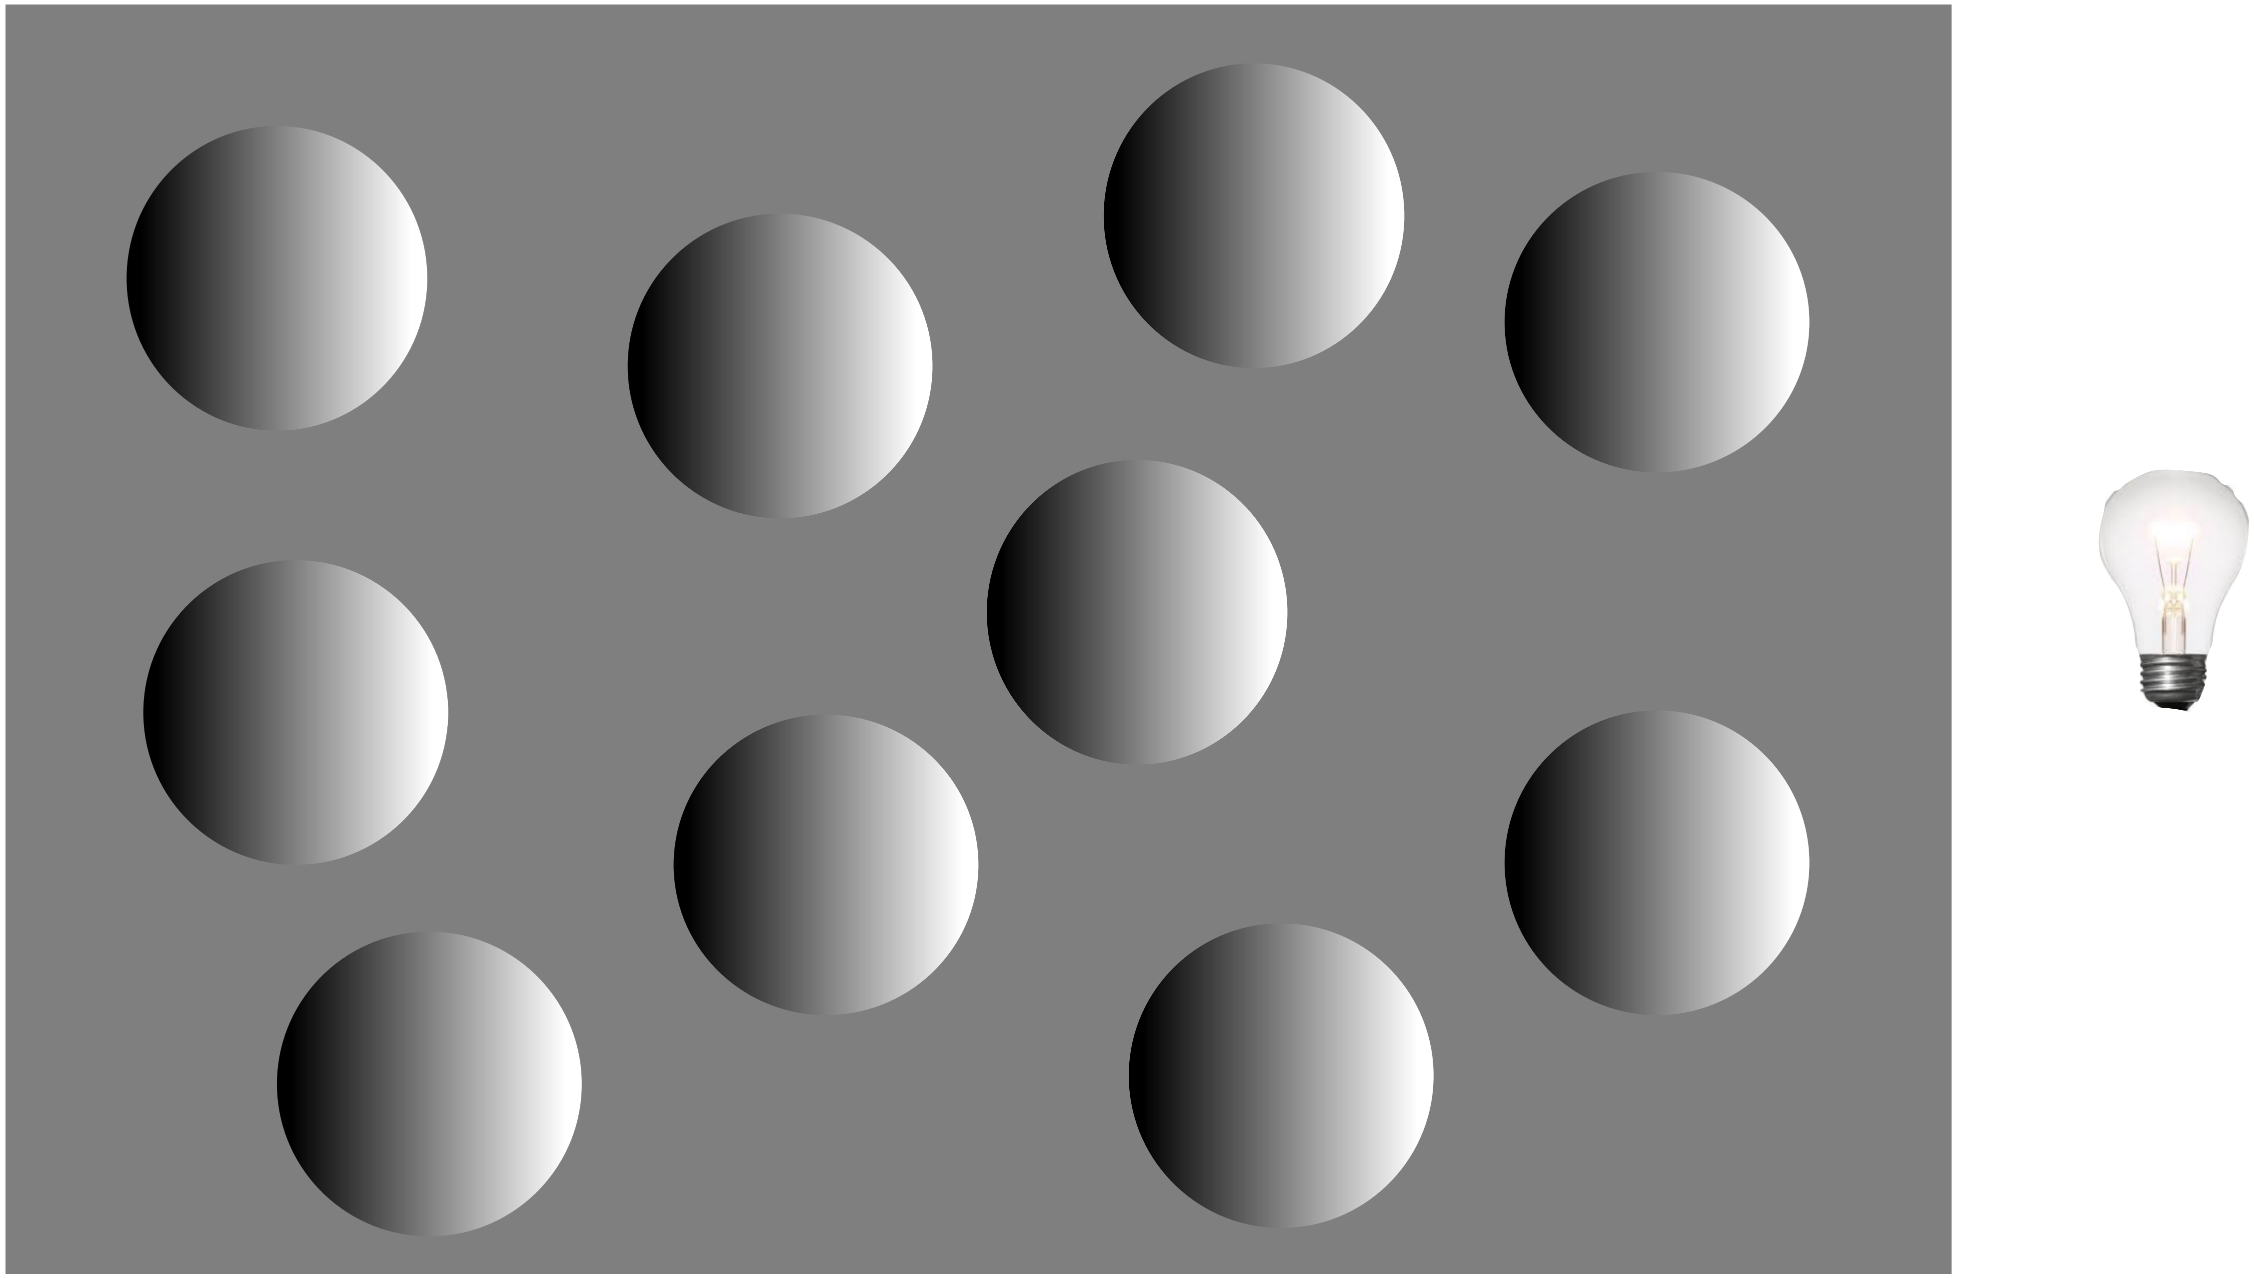

Solution 1: Recovery from a single image

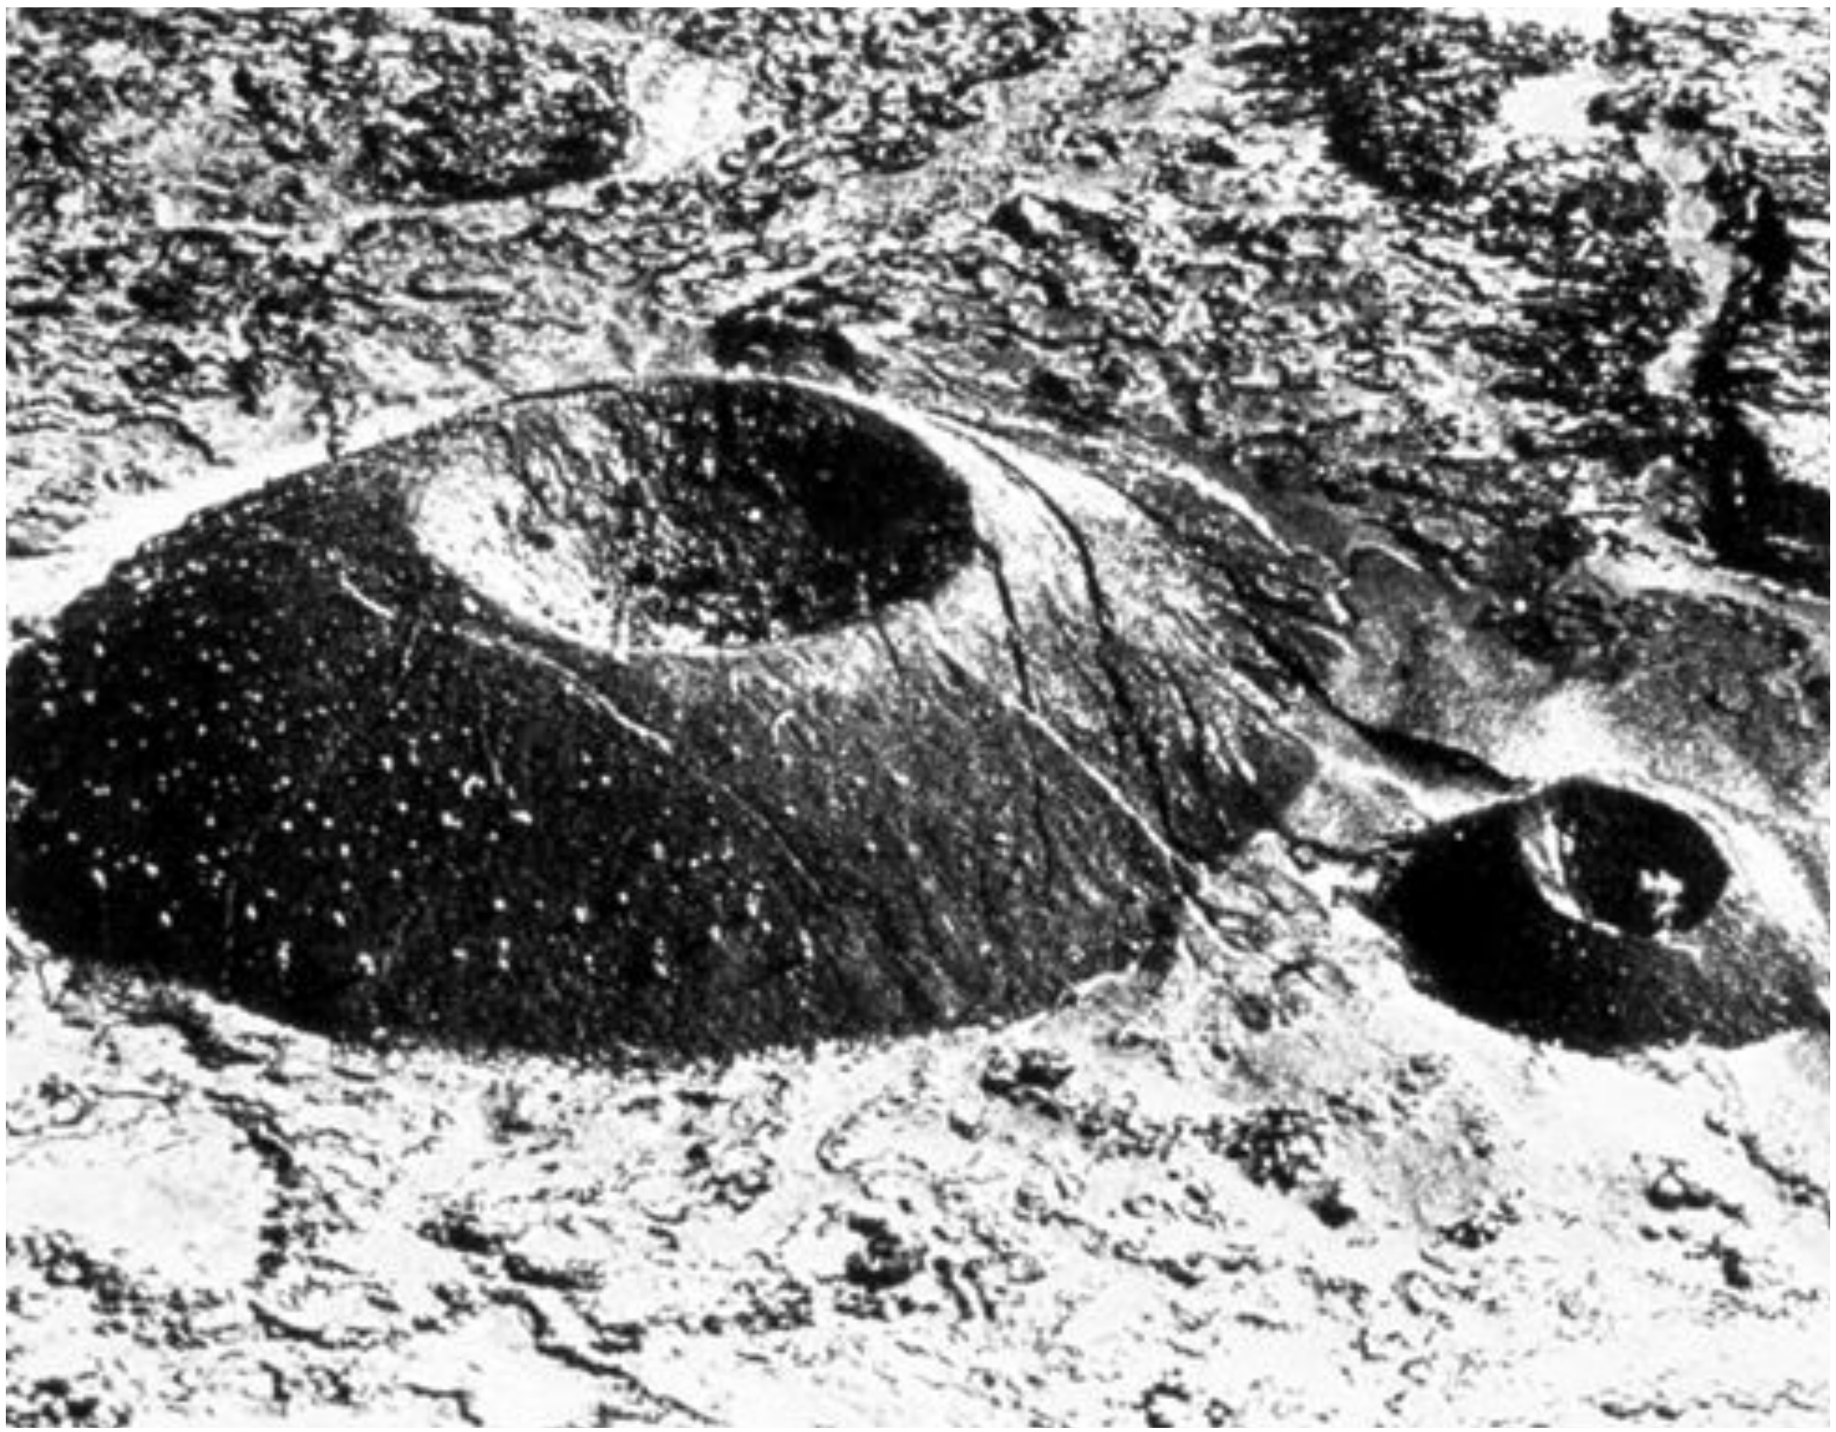

Solution 1: Recovery from a single image

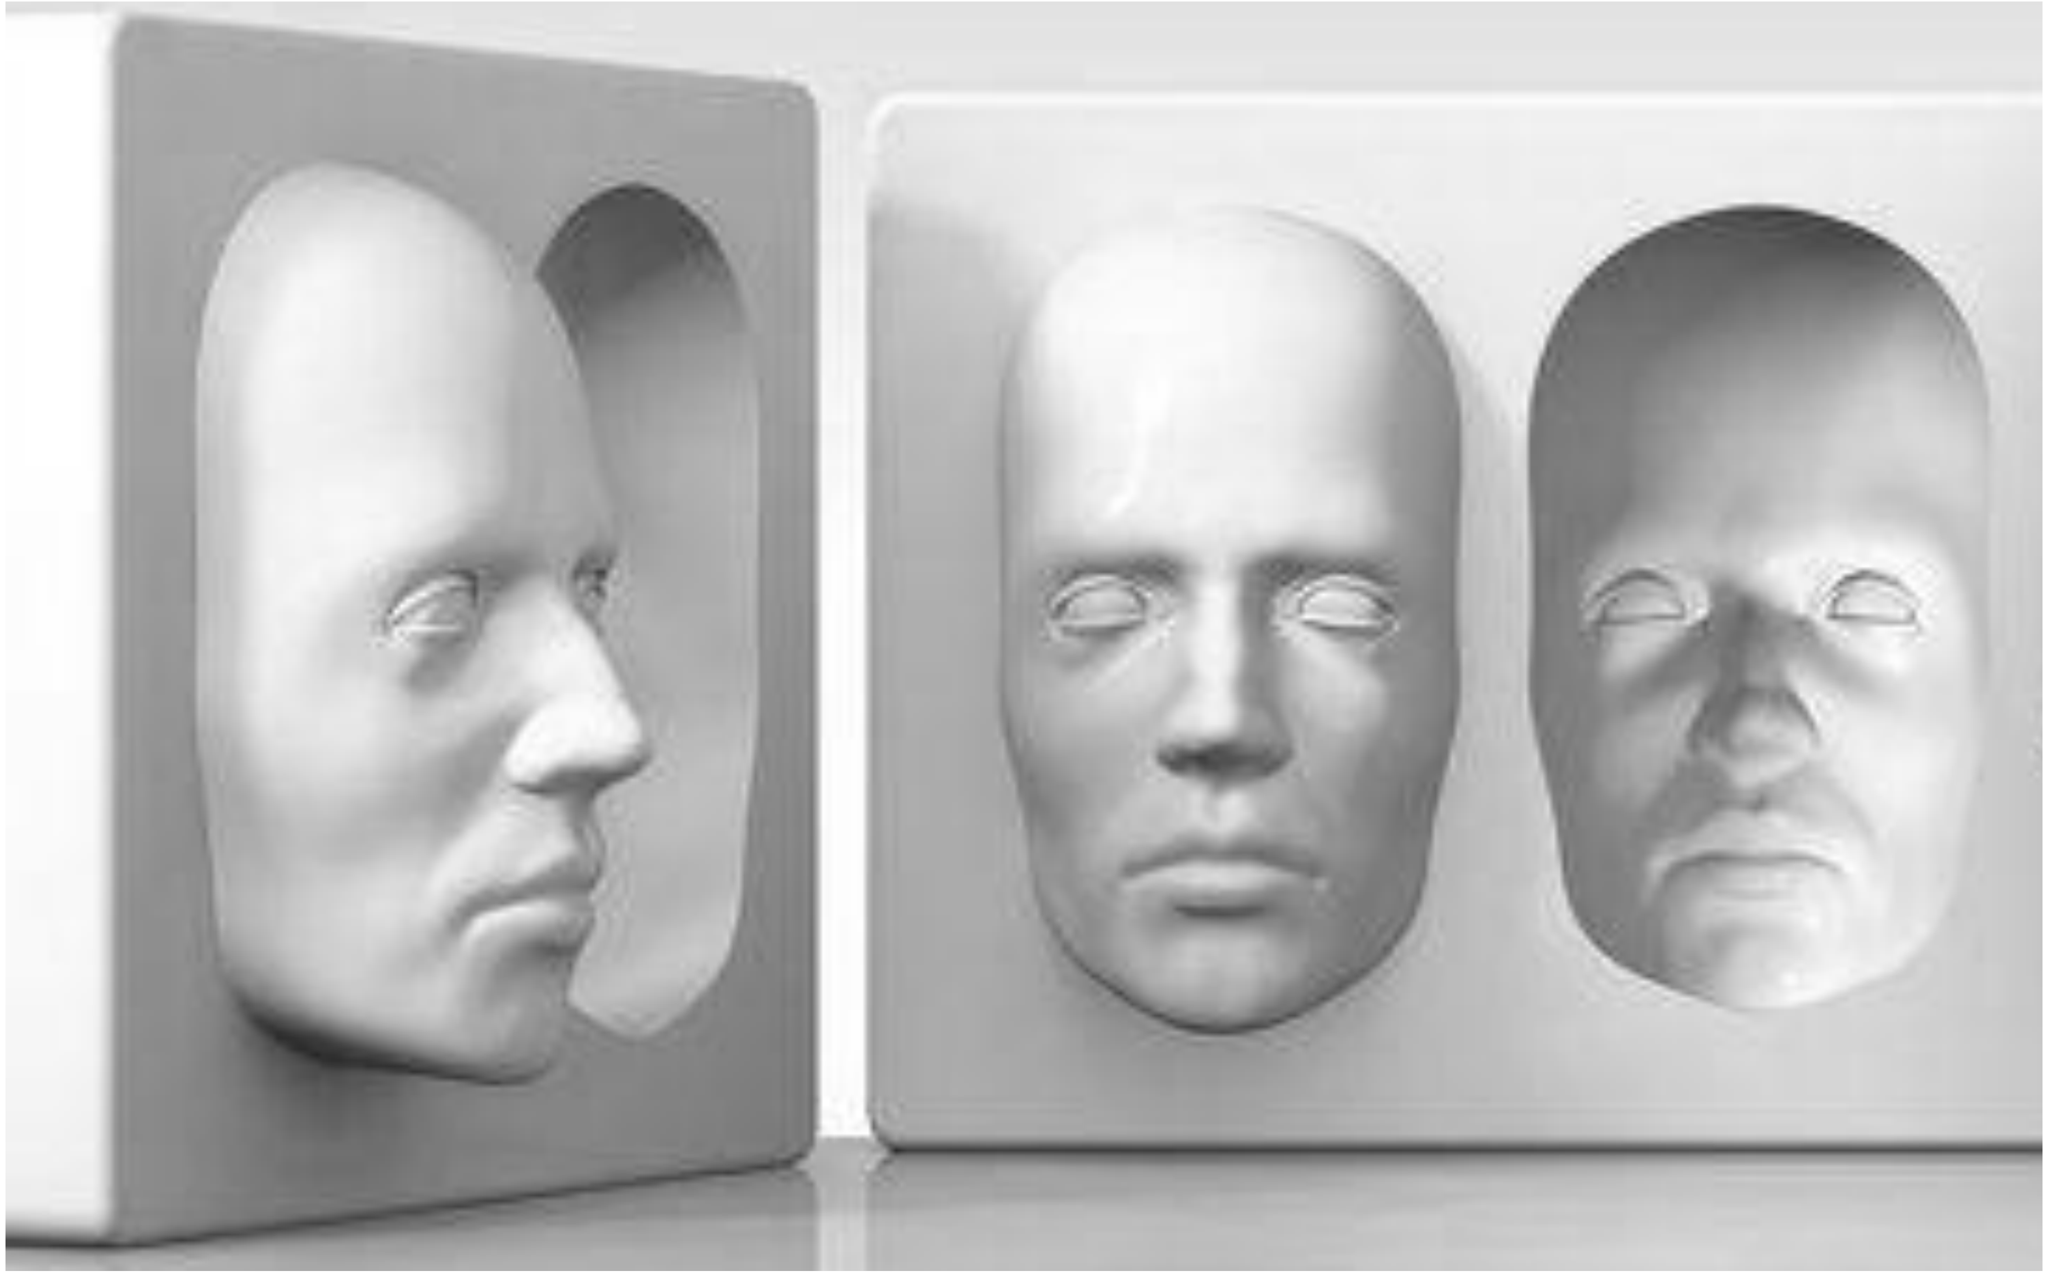

Solution 1: Recovery from a single image

- Step 1: Intrinsic image decomposition
  - Reflectance image $\rho(x, y)$
  - Shading image $L_i \mathbf{L} \cdot \mathbf{N}(x, y)$
  - Decomposition relies on priors on reflectance image
- What kind of priors?
  - Reflectance image captures the “paint” on an object surface
  - Surfaces tend to be of uniform color with sharp edges when color changes

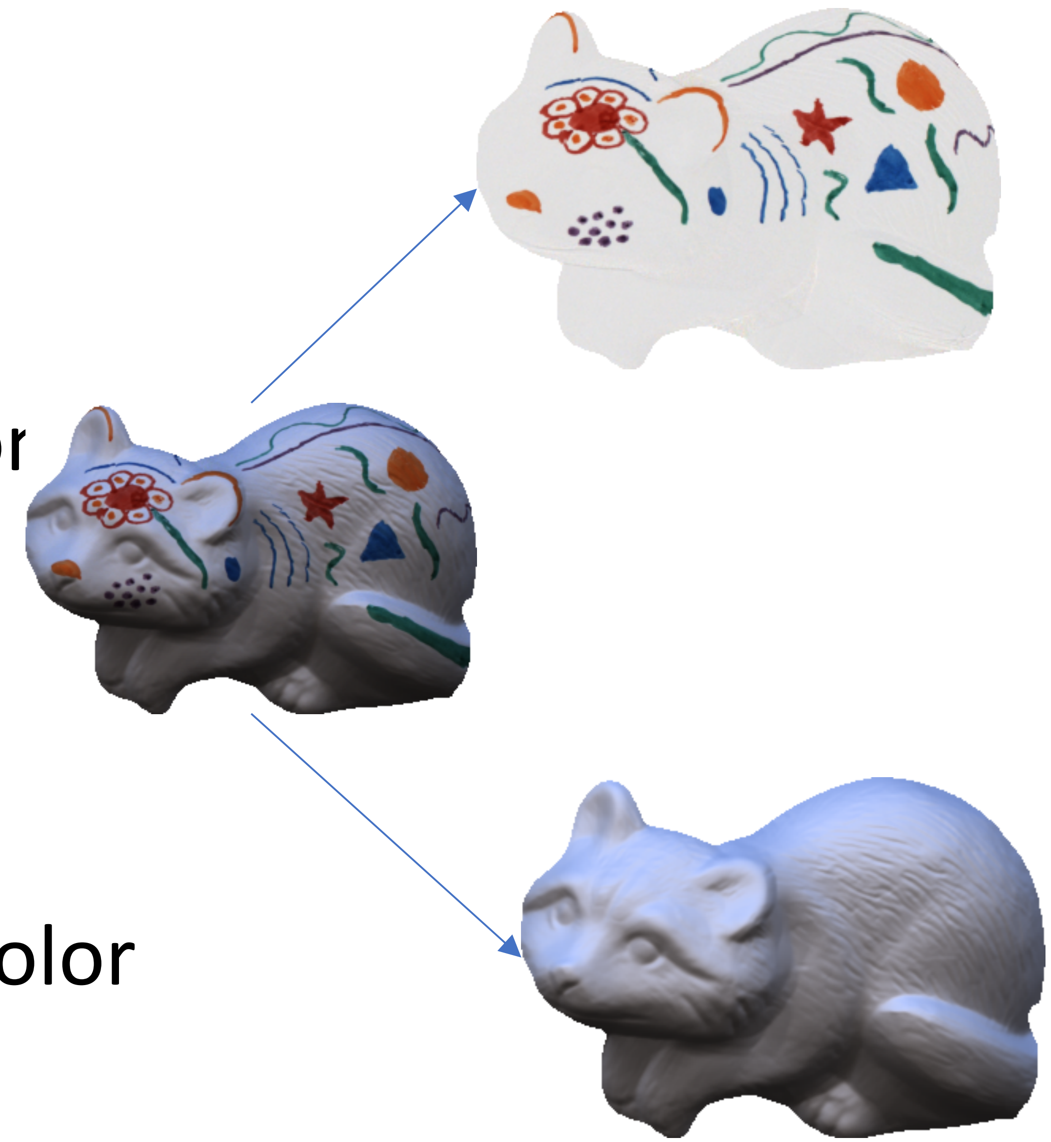

Images from Barron et al, TPAMI 13

Solution 1: Recovery from a single image

- Step 2: Decompose shading image into illumination and normals

  $$L_i \mathbf{L} \cdot \mathbf{N}(x, y)$$

  - Called Shape-From-Shading
  - Relies on priors on shape: shapes are smooth

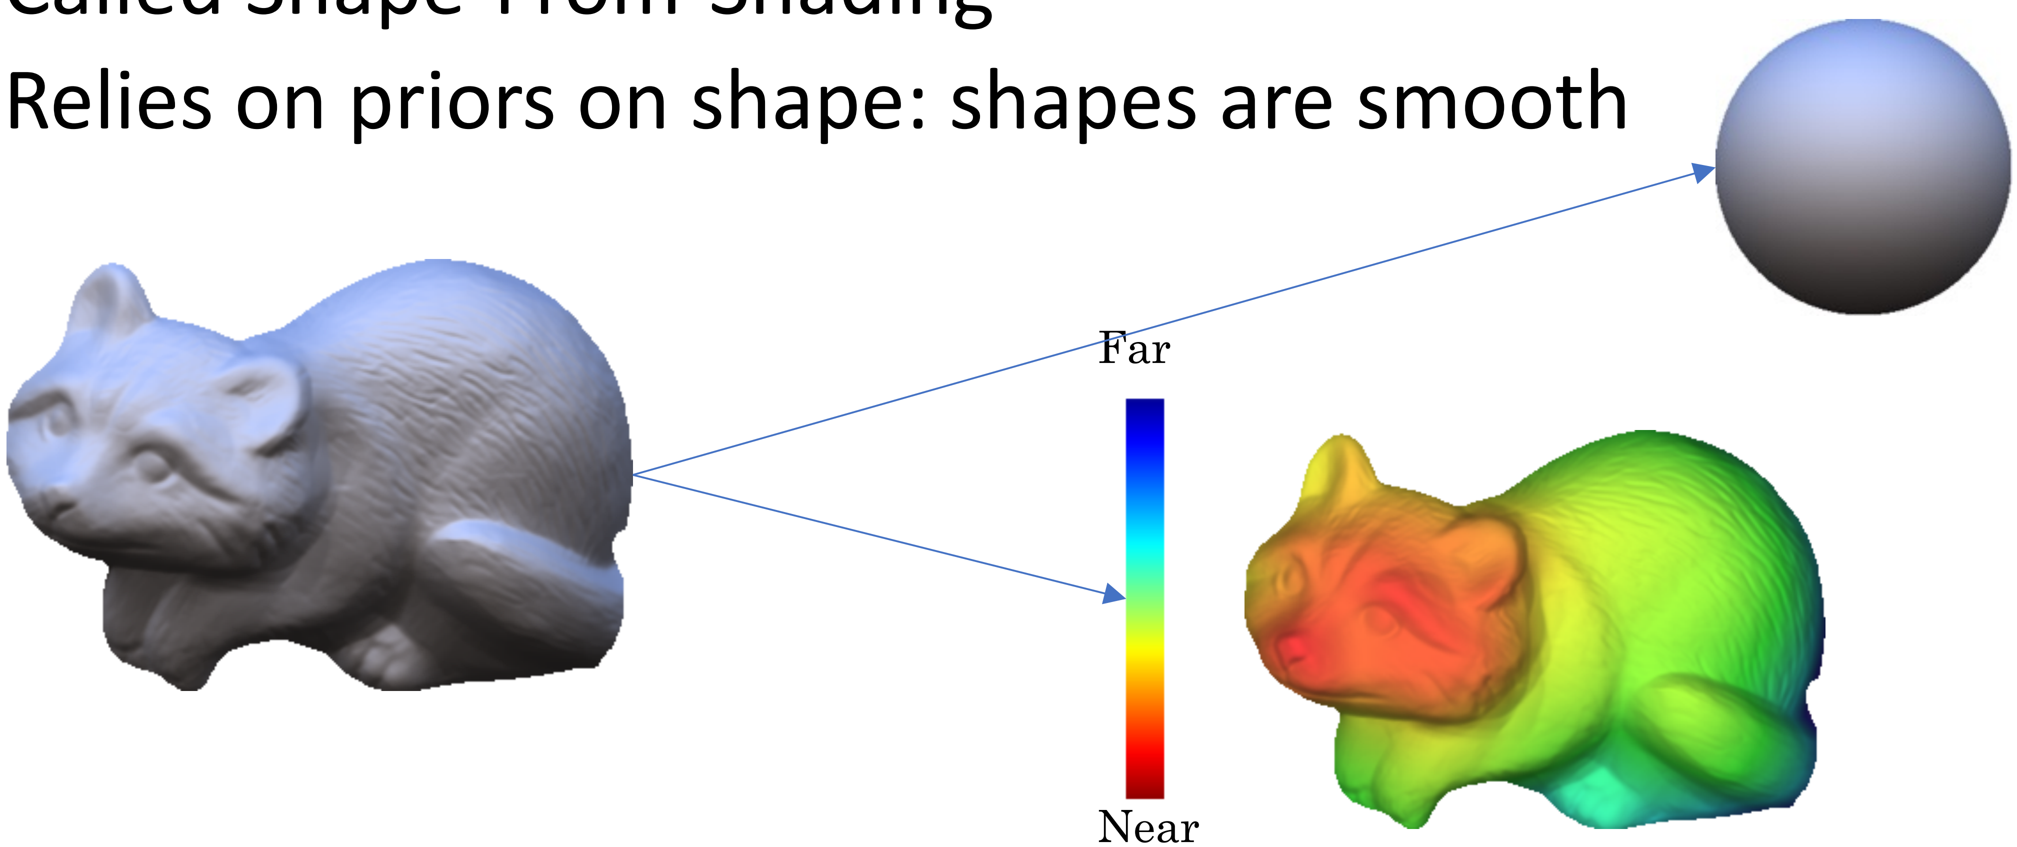

Photometric Stereo

Key Idea: use pixel brightness to understand shape

Solution 2: Recovery from multiple images

$$I(x, y) = \rho(x, y) L_i \mathbf{L} \cdot \mathbf{N}(x, y)$$

- Represents an equation in the albedo and normals
- Multiple images give constraints on albedo and normals
- Called *Photometric Stereo*

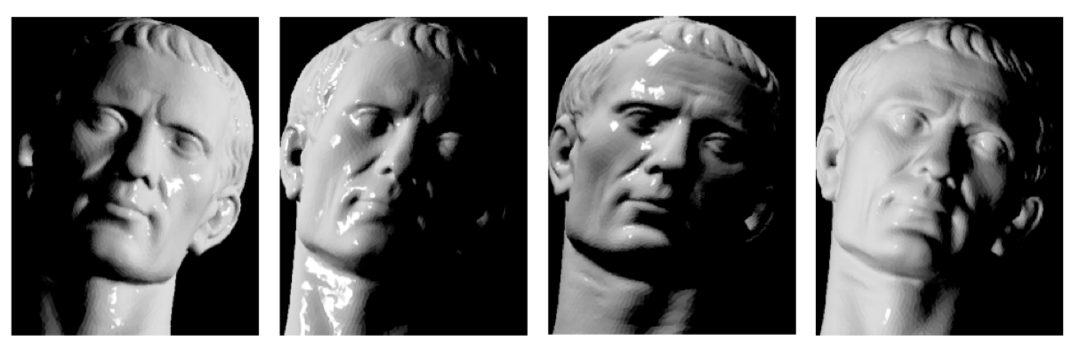

Image credit: Wikipedia

Photometric stereo - the math

$$I(x,y) = \rho(x,y) L_i \mathbf{L} \cdot \mathbf{N}(x,y)$$

- Consider single pixel
- Assume $L_i = 1$

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

$$I = \rho \mathbf{N}^T \mathbf{L}$$

- Write $\mathbf{G} = \rho \mathbf{N}$
- $\mathbf{G}$ is a 3-vector
  - Norm of **G** = ρ
  - Direction of **G** = **N**

Photometric stereo - the math

- Consider single pixel
- Assume $L_i = 1$

$$I = \rho \mathbf{N}^T \mathbf{L}$$

- Write $\mathbf{G} = \rho \mathbf{N}$
- **G** is a 3-vector
  - Norm of **G** = ρ
  - Direction of **G** = **N**

$$I = \mathbf{G}^T \mathbf{L} = \mathbf{L}^T \mathbf{G}$$

Photometric stereo - the math

$$I = \mathbf{L}^T \mathbf{G}$$

- Multiple images with different light sources but same viewing direction?

$$I_1 = \mathbf{L}_1^T \mathbf{G}$$

$$I_2 = \mathbf{L}_2^T \mathbf{G}$$

$$\vdots$$

$$I_k = \mathbf{L}_k^T \mathbf{G}$$

Photometric stereo - the math

$$I_1 = \mathbf{L}_1^T \mathbf{G}$$

$$I_2 = \mathbf{L}_2^T \mathbf{G}$$

$$\vdots$$

$$I_k = \mathbf{L}_k^T \mathbf{G}$$

- Assume lighting directions are known
- Each is a linear equation in G
- Stack everything up into a massive linear system of equations!

Photometric stereo - the math

$$I_1 = \mathbf{L}_1^T \mathbf{G}$$

$$I_2 = \mathbf{L}_2^T \mathbf{G}$$

$$\vdots$$

$$I_k = \mathbf{L}_k^T \mathbf{G}$$

k x 3 matrix of lighting directions

$$\mathbf{I} = \mathbf{L}^T \mathbf{G}$$

k x 1 vector of intensities

3x1 vector of unknowns

Photometric stereo - the math

$$\mathbf{I} = \mathbf{L}^T \mathbf{G}$$

k x 1 k x 3 3 x 1

$$\mathbf{G} = \mathbf{L}^{-T} \mathbf{I}$$

- What is the minimum value of k to allow recovery of G?
- How do we recover G if the problem is overconstrained?

Photometric stereo - the math

- How do we recover G if the problem is overconstrained?
  - More than 3 lights: more than 3 images

- Least squares

$$\min_{\mathbf{G}} \|\mathbf{I} - \mathbf{L}^T\mathbf{G}\|^2$$

- Solved using normal equations

$$\mathbf{G} = (\mathbf{L}\mathbf{L}^T)^{-1}\mathbf{L}\mathbf{I}$$

Normal equations

$$\|\mathbf{I} - \mathbf{L}^T\mathbf{G}\|^2 = \mathbf{I}^T\mathbf{I} + \mathbf{G}^T\mathbf{L}\mathbf{L}^T\mathbf{G} - 2\mathbf{G}^T\mathbf{L}\mathbf{I}$$

- Take derivative with respect to **G** and set to 0

$$2\mathbf{L}\mathbf{L}^T\mathbf{G} - 2\mathbf{L}\mathbf{I} = 0$$

$$\Rightarrow \mathbf{G} = (\mathbf{L}\mathbf{L}^T)^{-1}\mathbf{L}\mathbf{I}$$

Estimating normals and albedo from G

- Recall that $\mathbf{G} = \rho \mathbf{N}$

$$\|\mathbf{G}\| = \rho$$

$$\frac{\mathbf{G}}{\|\mathbf{G}\|} = \mathbf{N}$$

Multiple pixels

- We've looked at a single pixel till now
- How do we handle multiple pixels?
- Essentially independent equations!

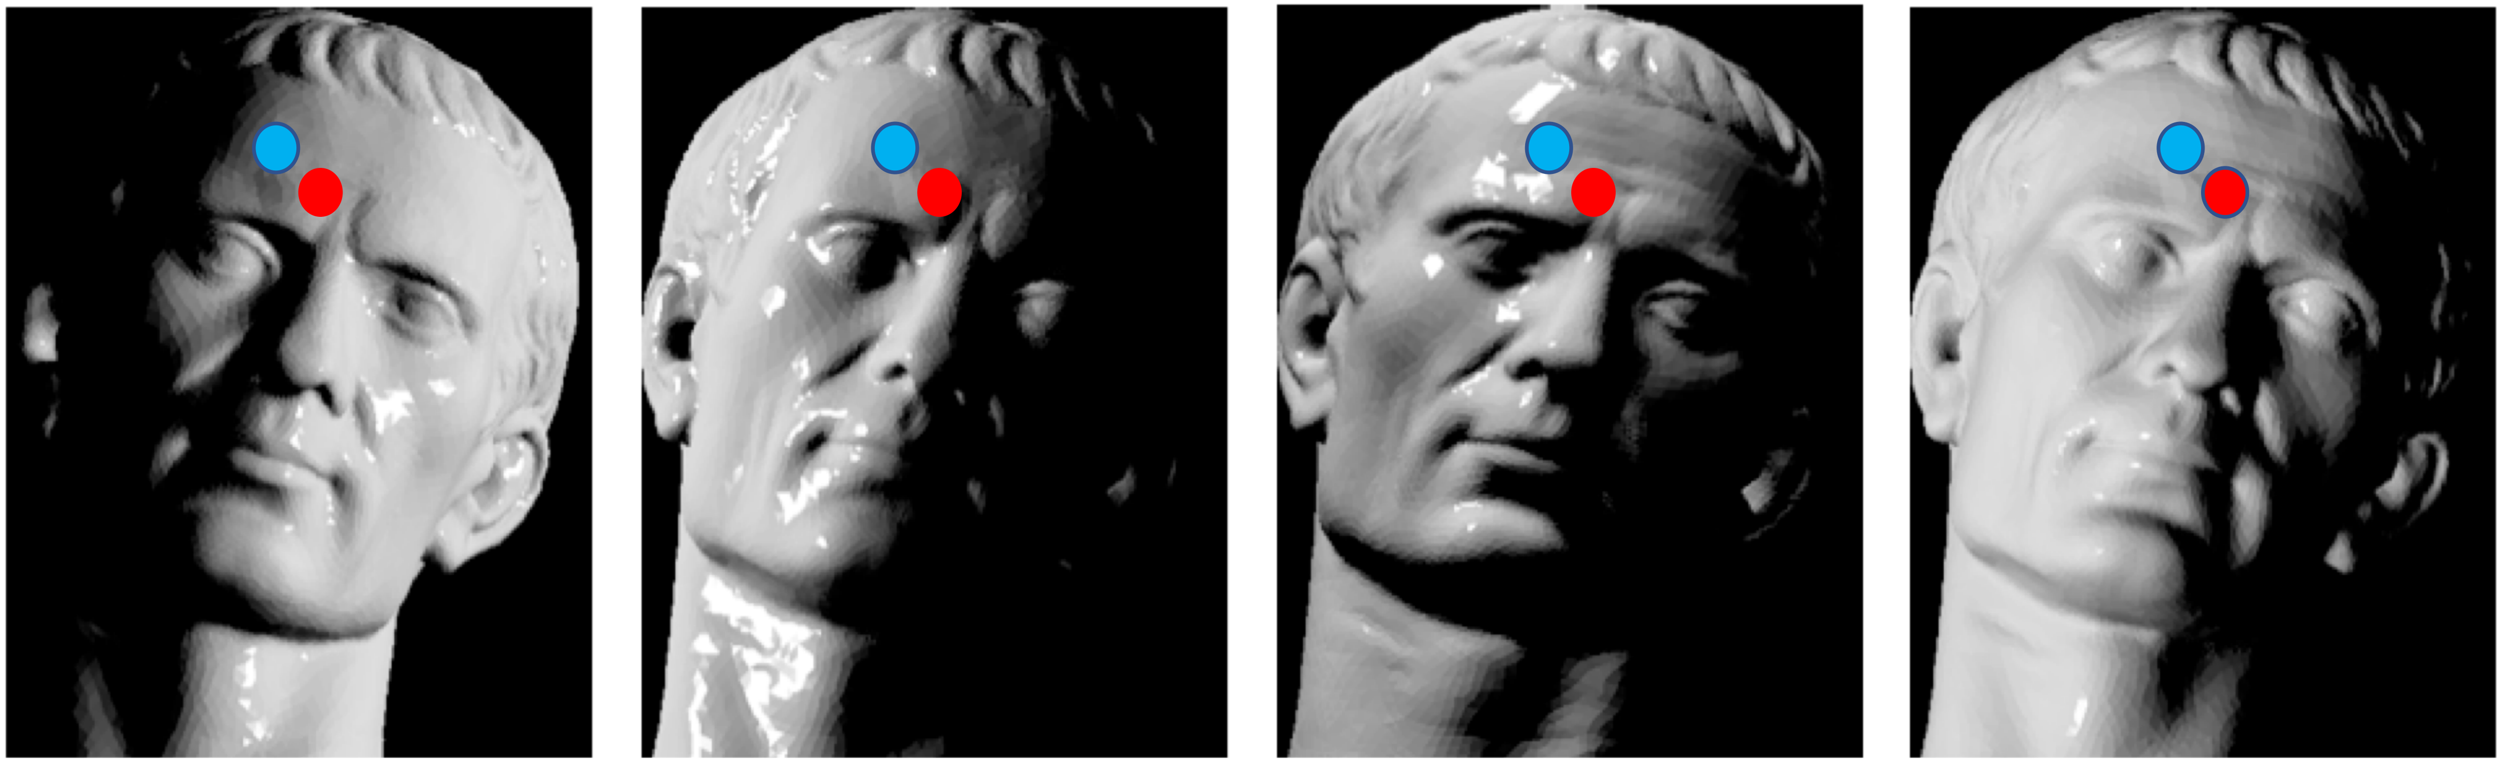

Multiple pixels: matrix form

- Note that all pixels share the same set of lights

$$\mathbf{I}^{(1)} = \mathbf{L}^T \mathbf{G}^{(1)}$$

$$\mathbf{I}^{(2)} = \mathbf{L}^T \mathbf{G}^{(2)}$$

$$\vdots$$

$$\mathbf{I}^{(n)} = \mathbf{L}^T \mathbf{G}^{(n)}$$

Multiple pixels: matrix form

- Can stack these into *columns* of a matrix

$$\mathbf{I}^{(1)} = \mathbf{L}^T \mathbf{G}^{(1)}$$

$$\mathbf{I}^{(2)} = \mathbf{L}^T \mathbf{G}^{(2)}$$

$$\vdots$$

$$\mathbf{I}^{(n)} = \mathbf{L}^T \mathbf{G}^{(n)}$$

$$\begin{bmatrix}\mathbf{I}^{(1)} & \mathbf{I}^{(2)} & \cdots & \mathbf{I}^{(n)}\end{bmatrix} = \mathbf{L}^T \begin{bmatrix}\mathbf{G}^{(1)} & \mathbf{G}^{(2)} & \cdots & \mathbf{G}^{(n)}\end{bmatrix}$$

$$\boxed{\mathbf{I} = \mathbf{L}^T \mathbf{G}}$$

Multiple pixels: matrix form

$$\mathbf{I} = \mathbf{L}^T \mathbf{G}$$

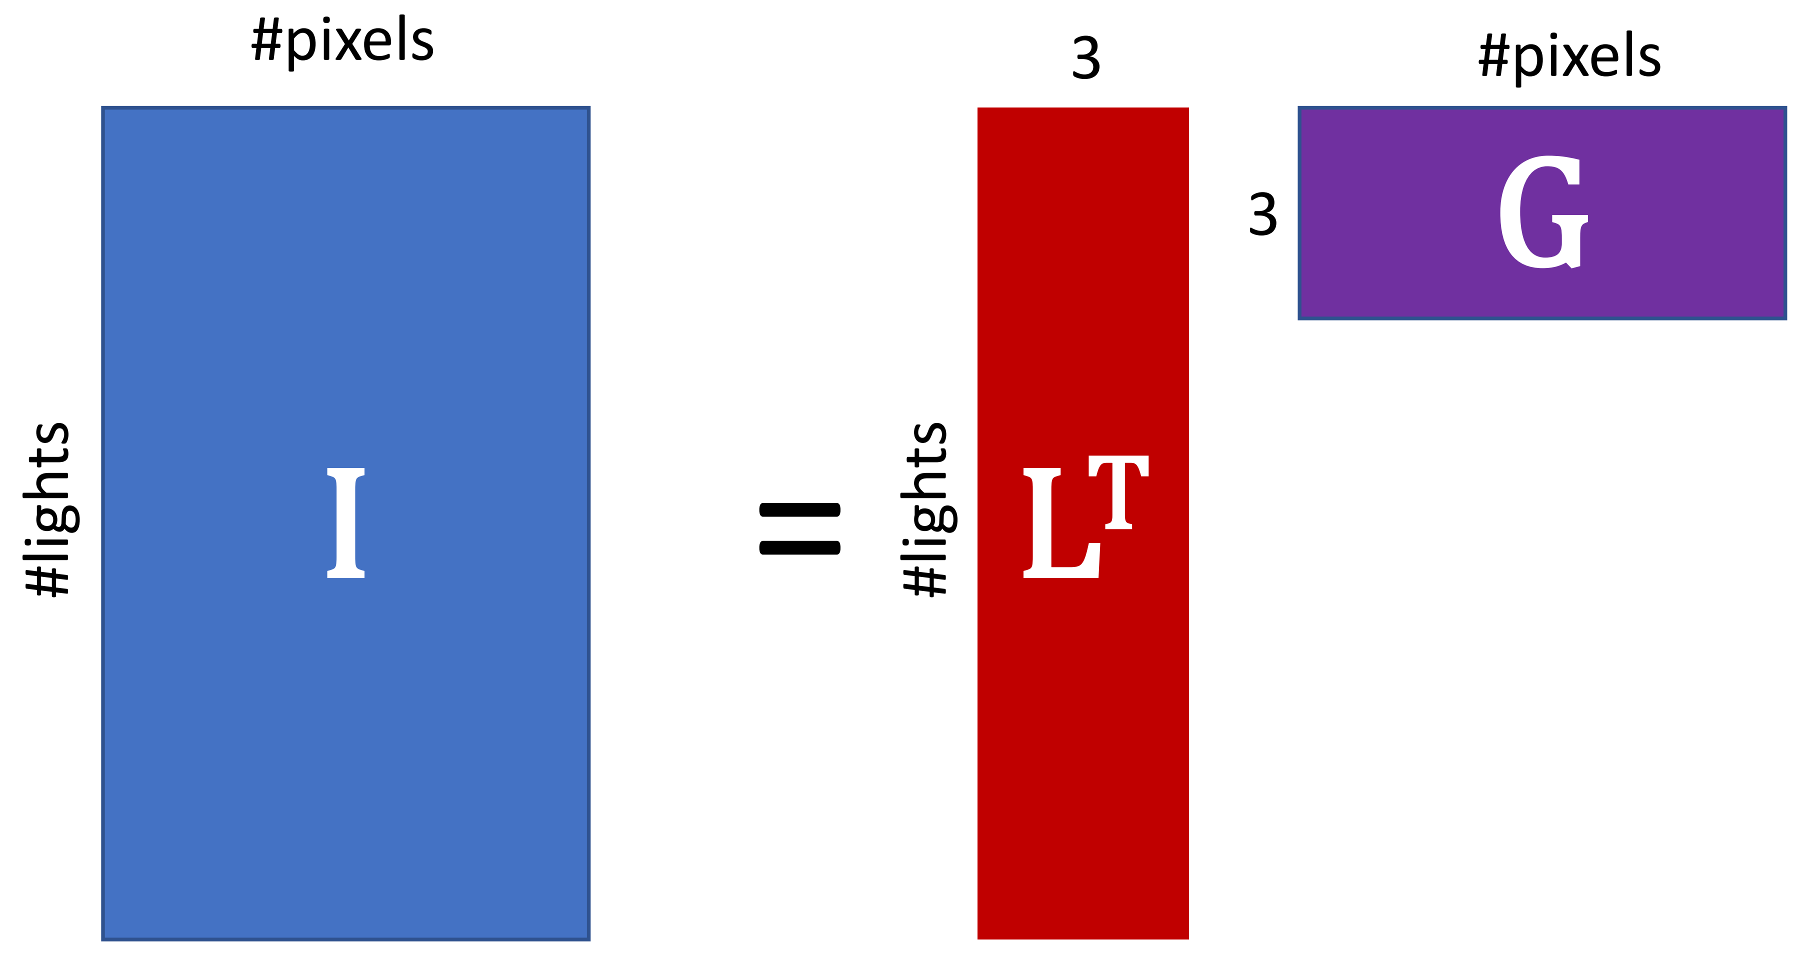

Determining Light Directions

- Trick: Place a mirror ball in the scene.

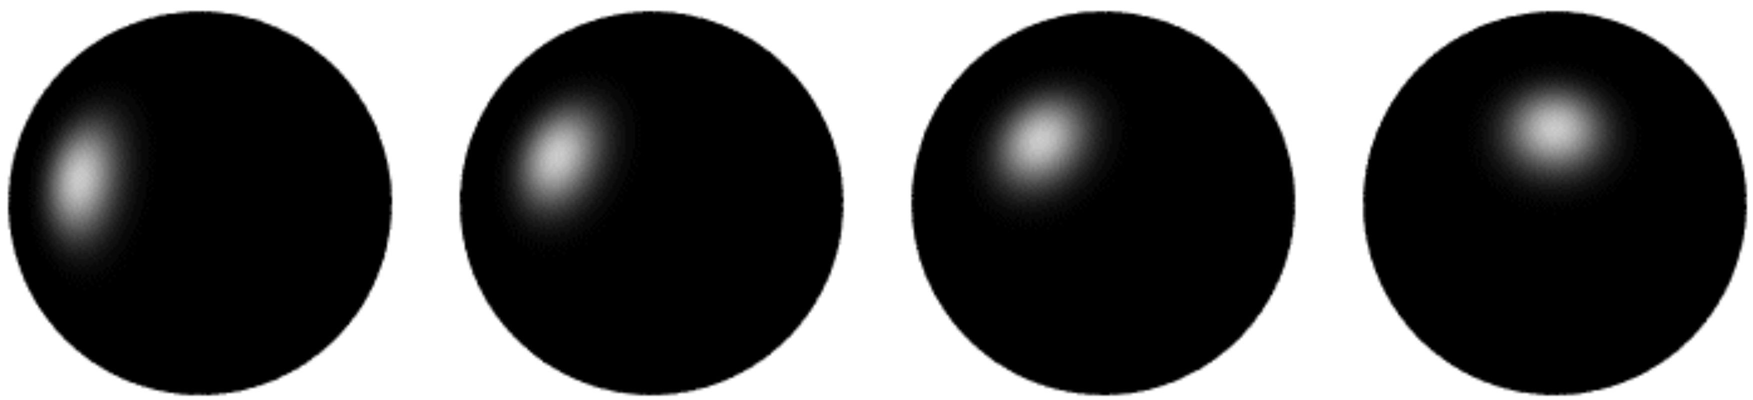

- The location of the highlight is determined by the light source direction.

Real-World HDR Lighting Environments

Funston Beach

Eucalyptus Grove

Uffizi Gallery

Grace Cathedral

Lighting Environments from the Light Probe Image Gallery:
http://www.debevec.org/Probes/

Mirrored Sphere


Canon
Canon
REMOTE SWITCH

Extreme HDR Image Series

1 sec
f/4

1/4 sec
f/4

1/30 sec
f/4

1/30 sec
f/16

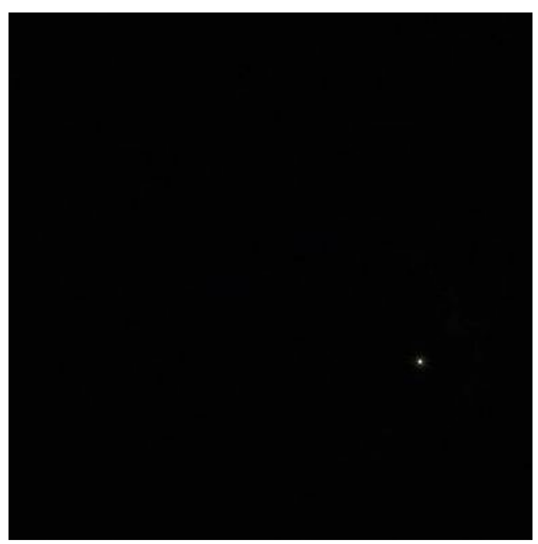

1/250 sec
f/16

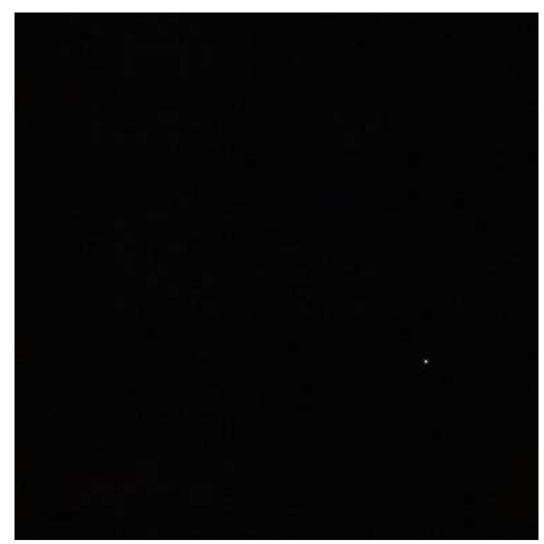

1/1000 sec
f/16

1/8000 sec f/16

Extreme HDR Image Series

sun closeup

1 sec
f/4

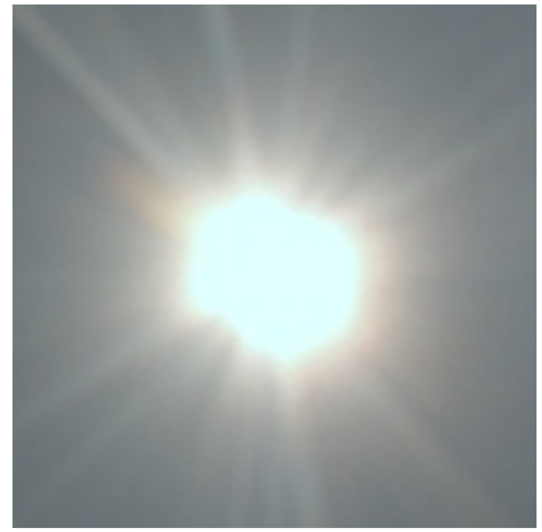

1/4 sec
f/4

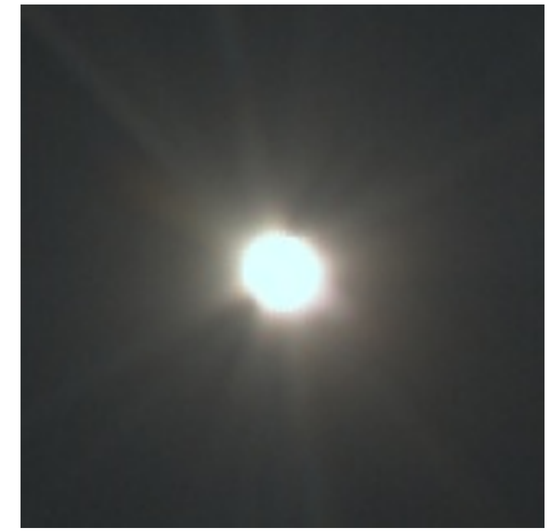

1/30 sec
f/4

1/30 sec
f/16

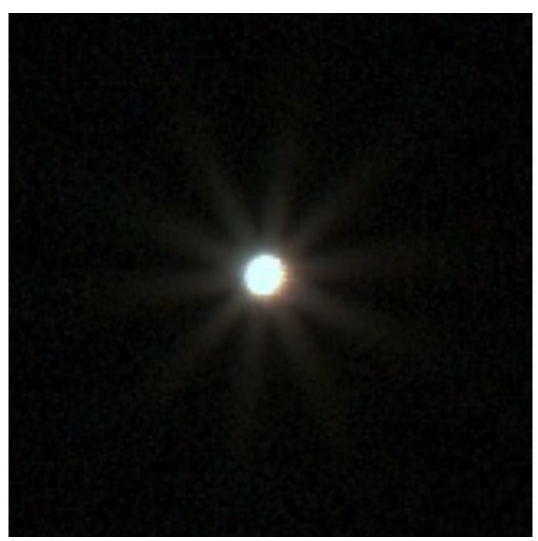

1/250 sec
f/16

1/1000 sec
f/16

1/8000 sec f/16
only image that does not saturate!

HDRI Sky Probe

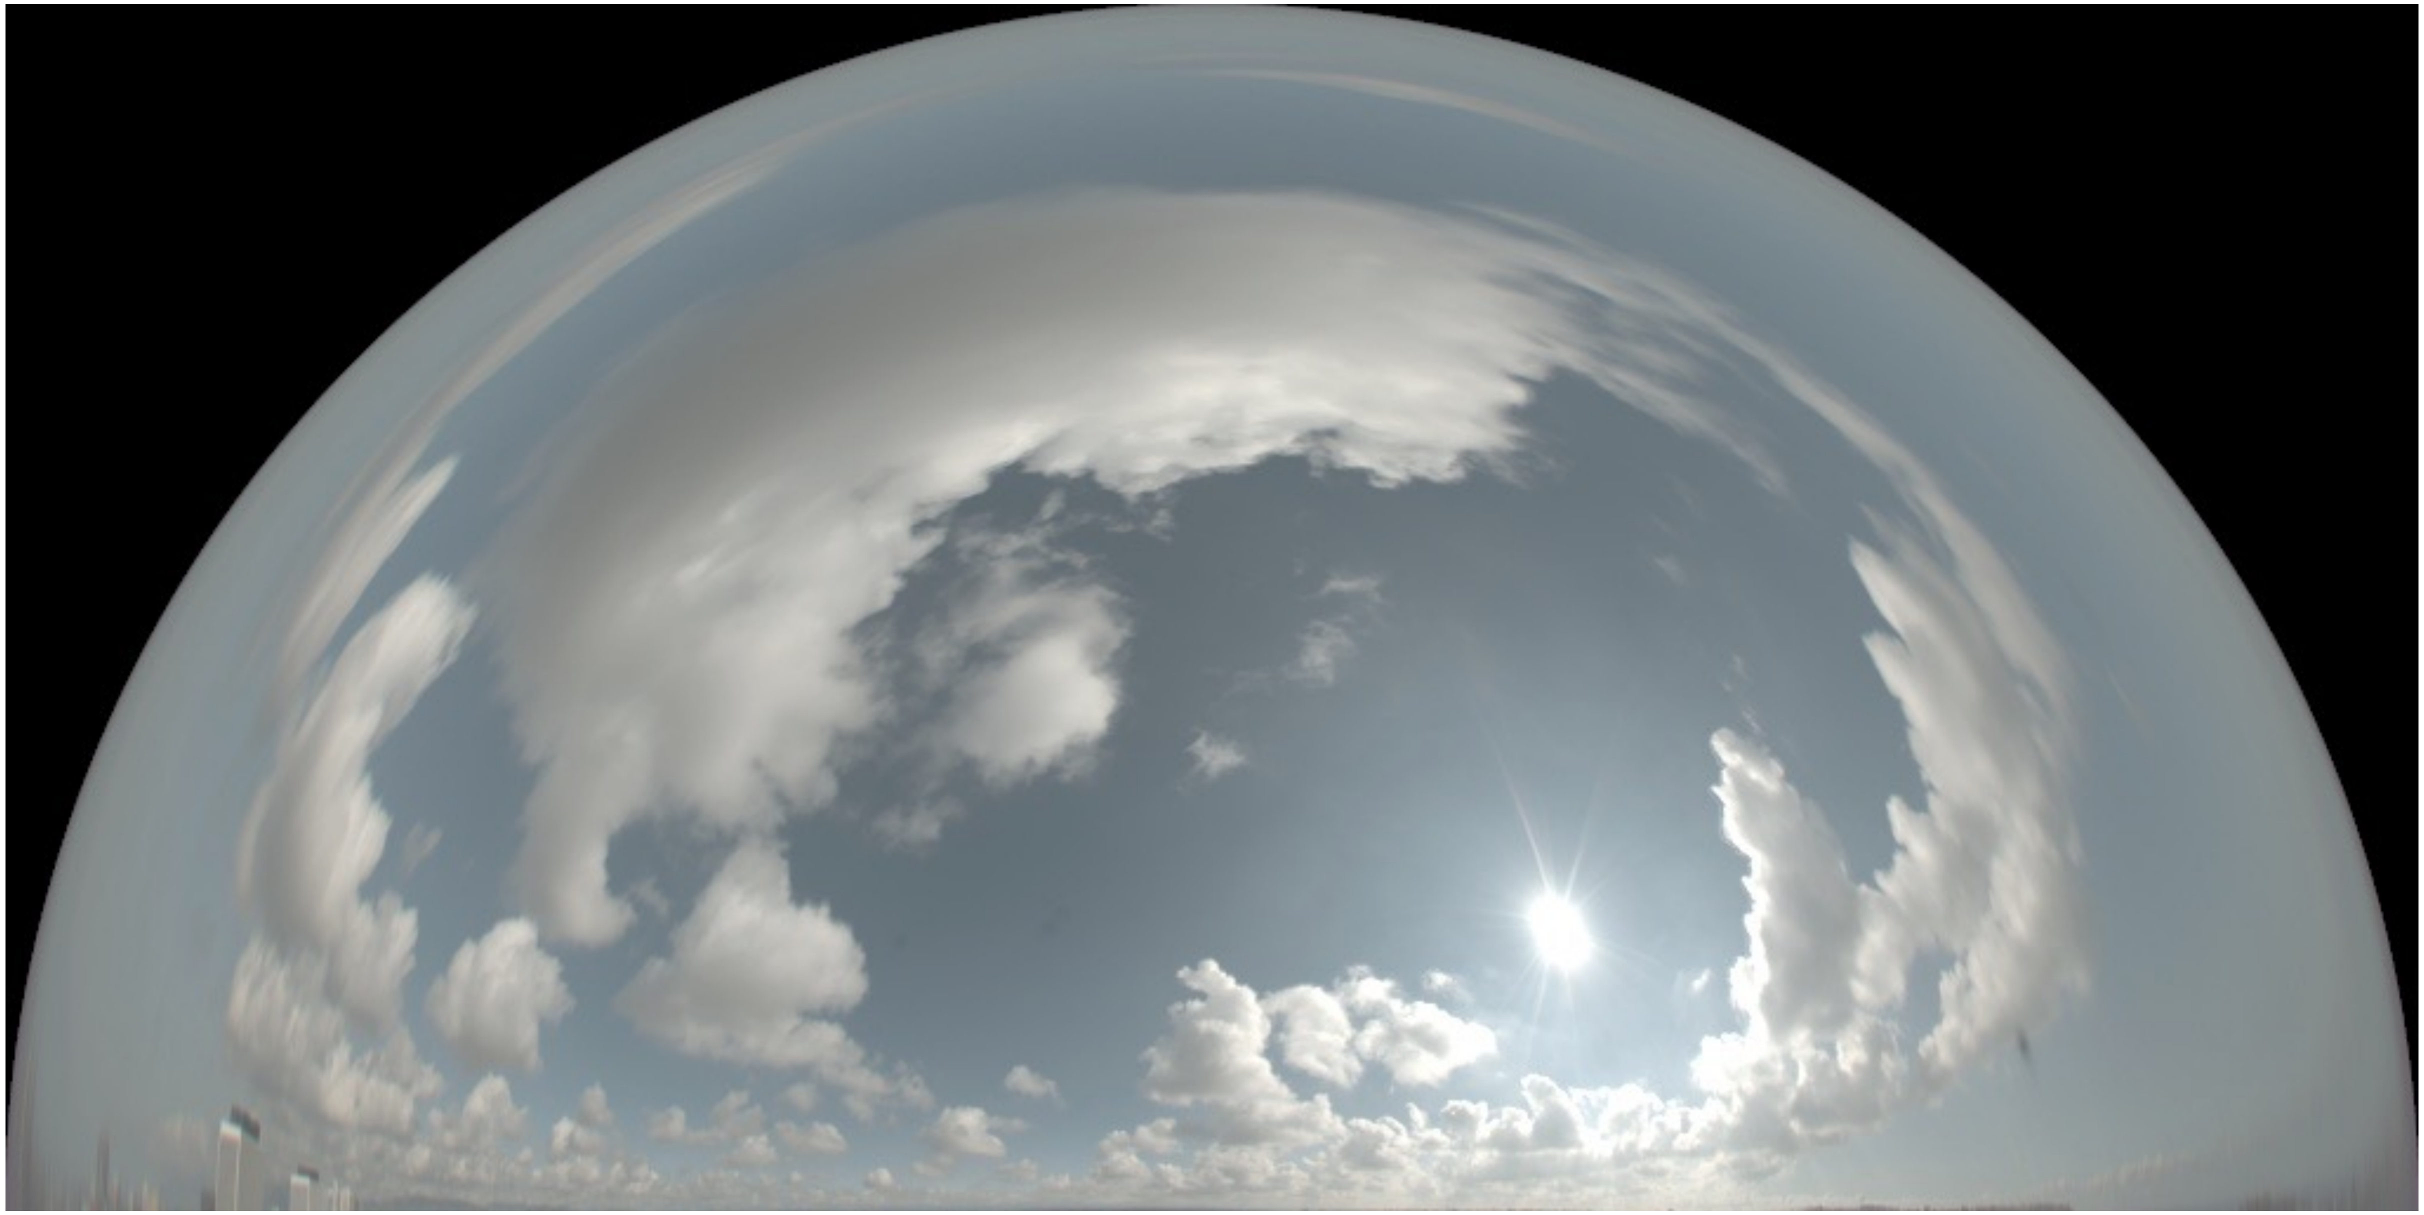

Determining Light Directions

- For a perfect mirror, the light is reflected across N:

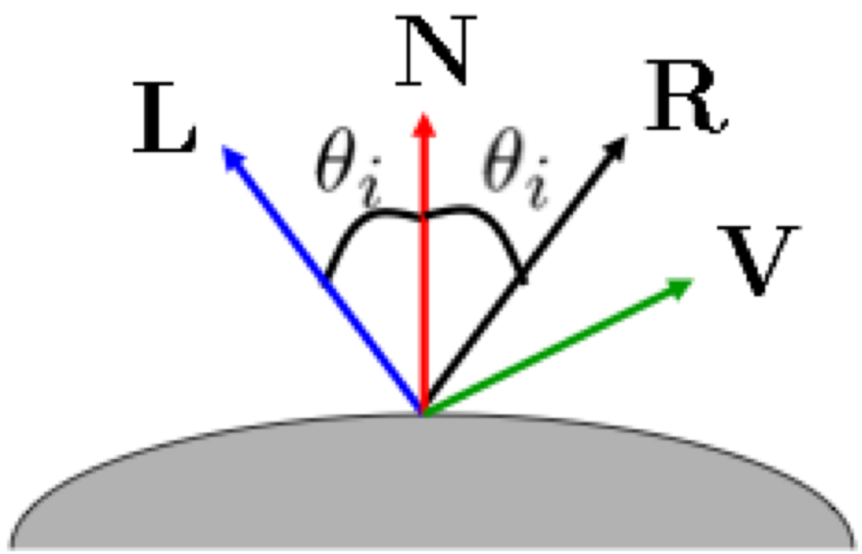

$$I_e = \begin{cases} I_i & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

Determining Light Directions

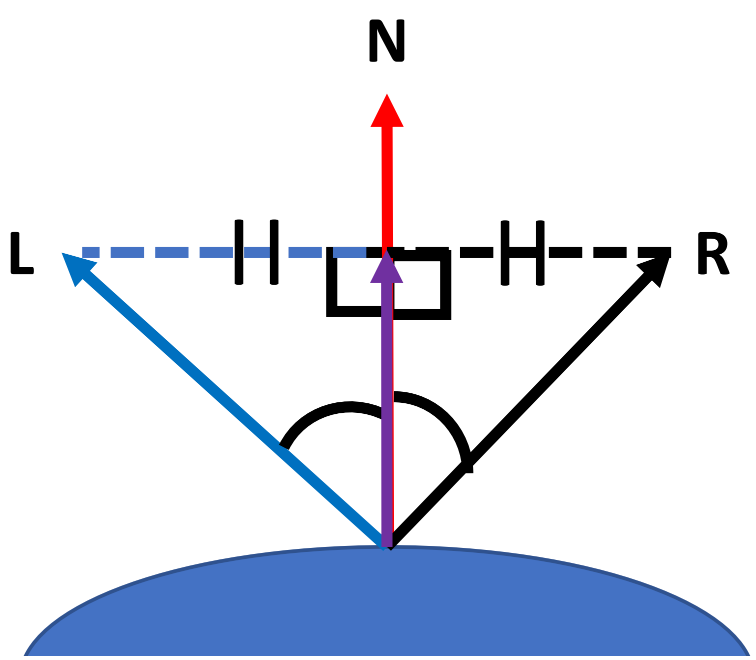

$$= (N \cdot R)N$$

$$= R - (N \cdot R)N$$

$$= R - (N \cdot R)N$$

$$= \quad - 2$$

$$= R - 2(R - N \cdot R)N$$

$$= 2(N \cdot R)N - R$$

So the light source direction is given by:

$$L = 2(N \cdot R)N - R$$

Photometric Stereo

What results can you get?

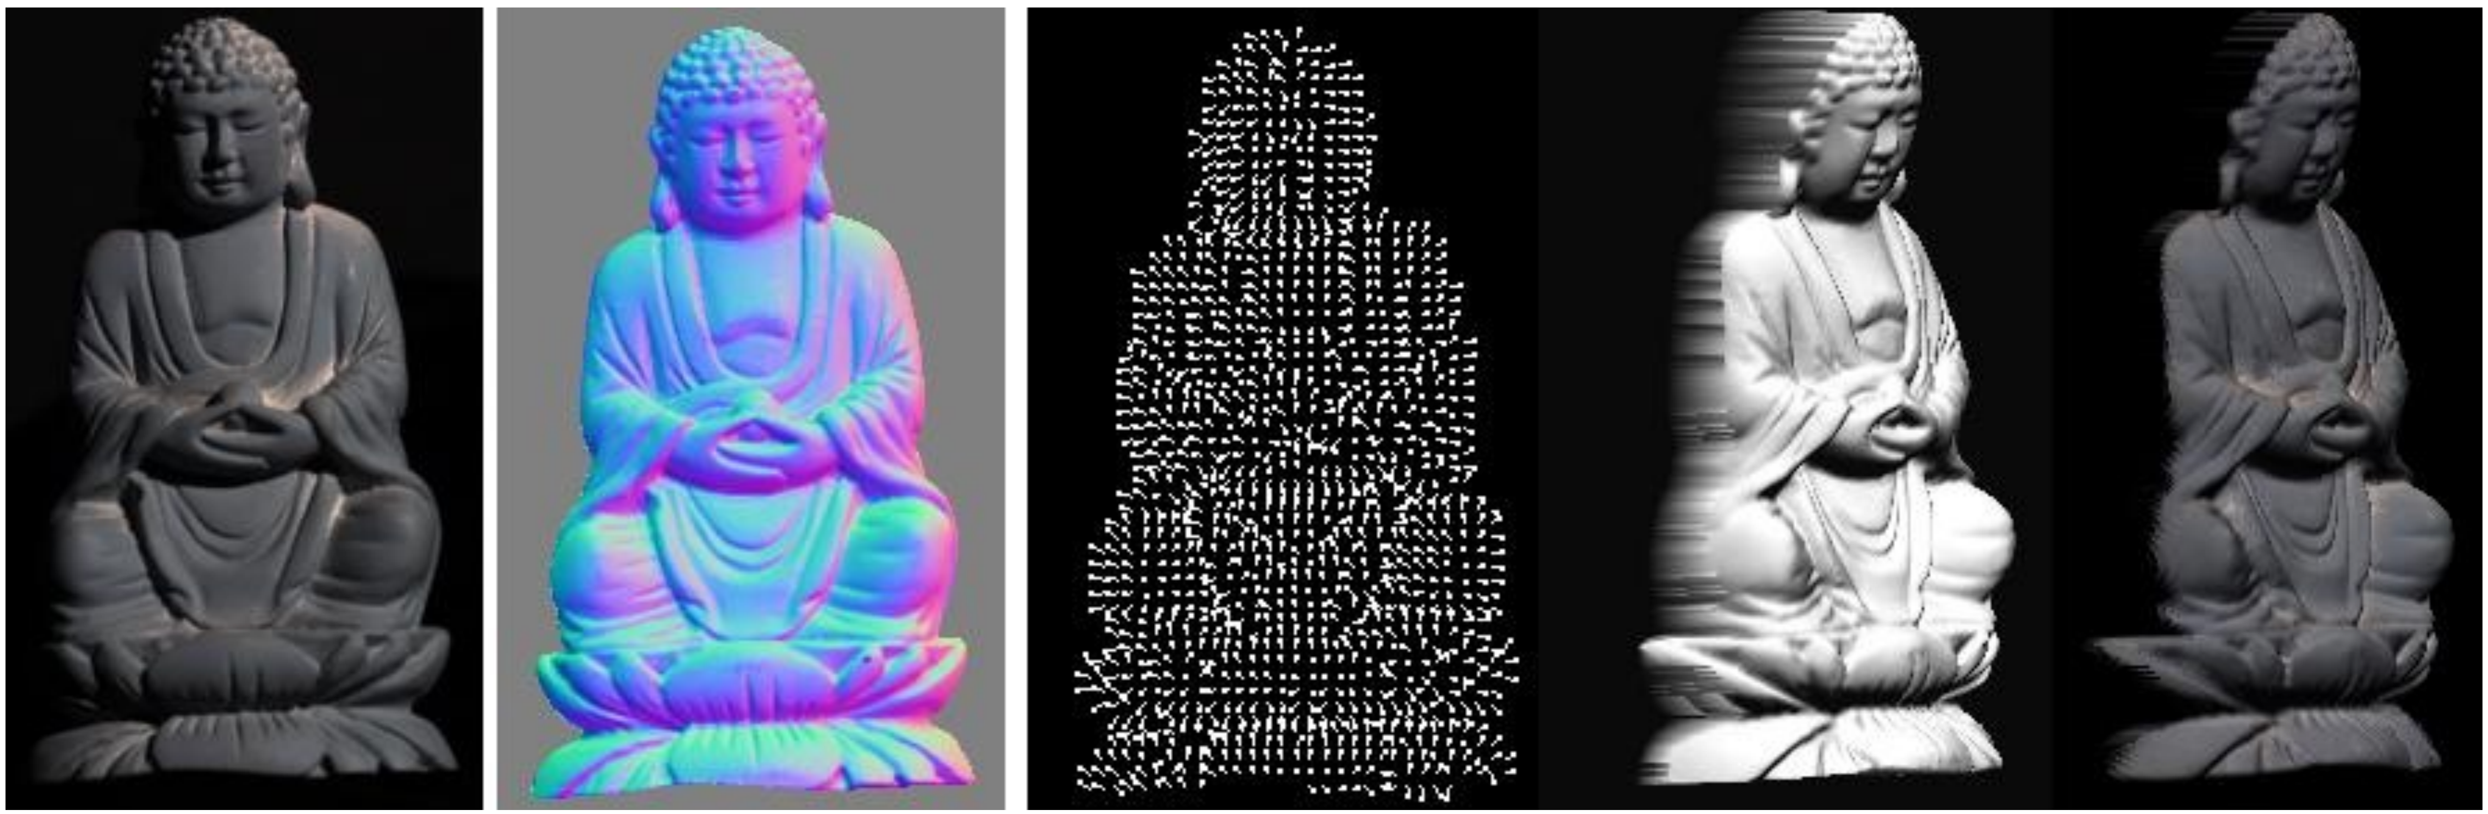

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Results

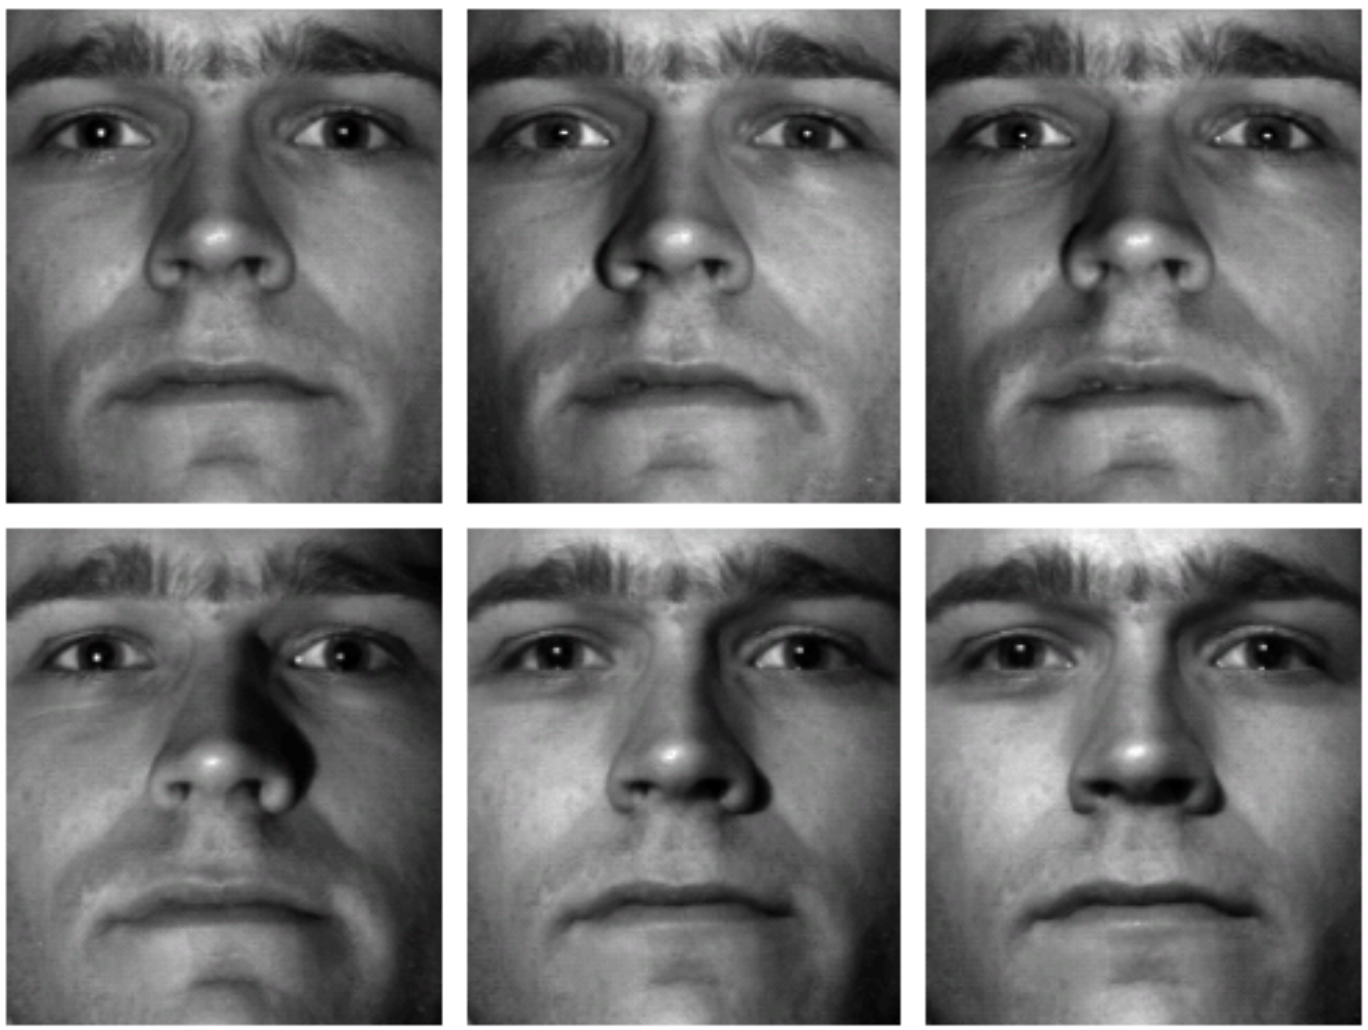

from Athos Georghiades

Unknown Lighting

Surface normals Light directions

$$I = kN \cdot \ell L$$

Diffuse albedo

Light intensity

Unknown Lighting

Surface normals, scaled by albedo

Light directions, scaled by intensity

$$I = N \cdot L$$

Unknown Lighting

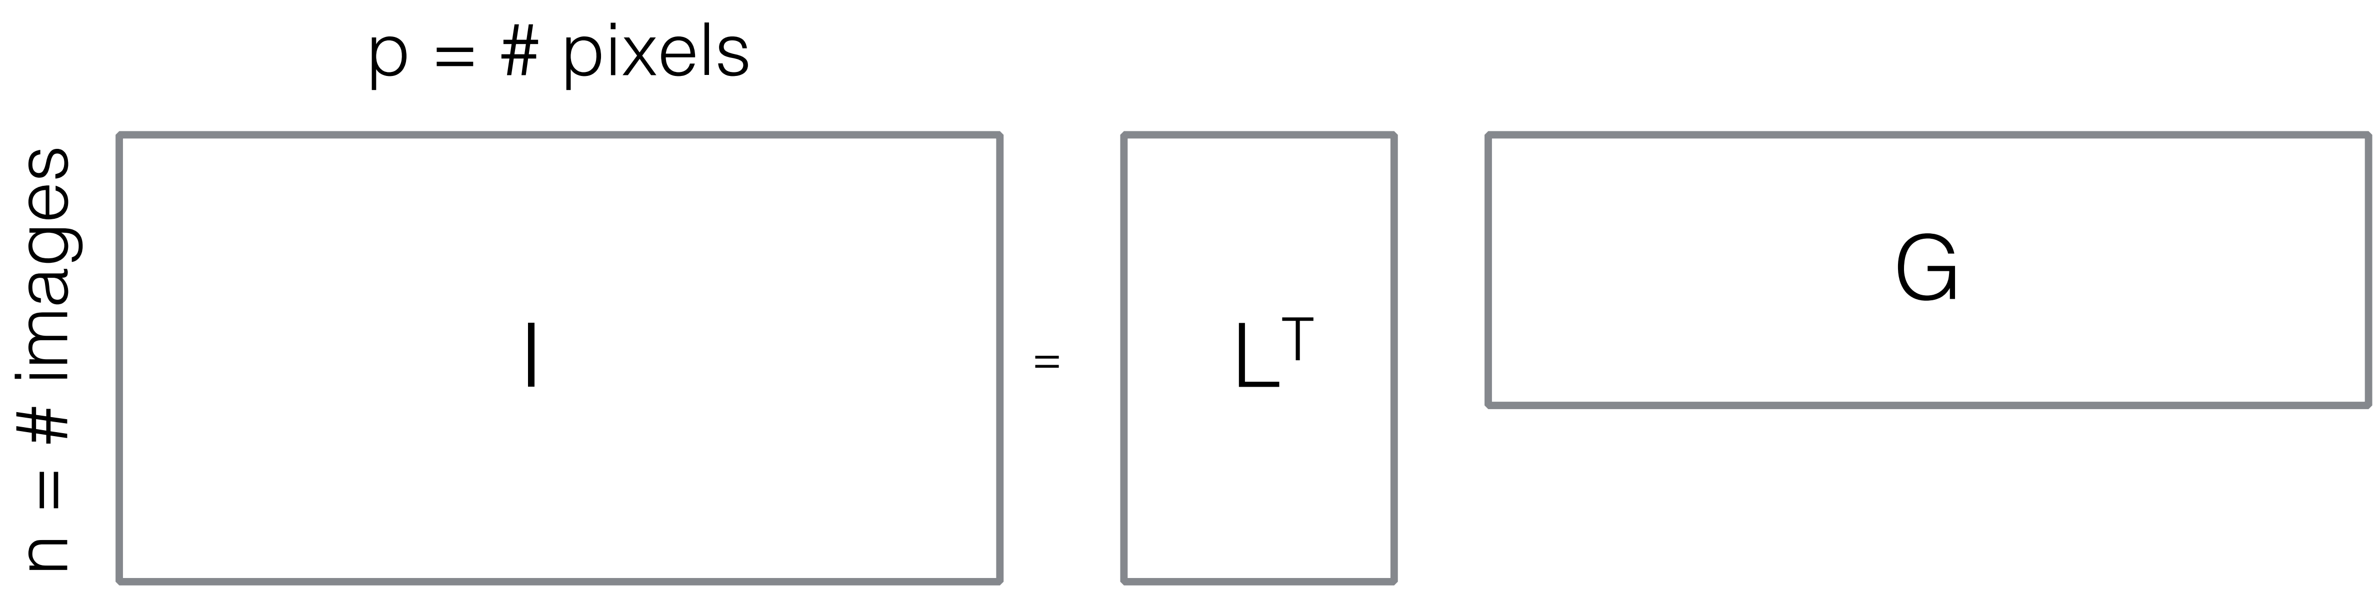

Unknown Lighting

Measurements (one image per row)

Light directions (scaled by intensity)

Surface normals (scaled by albedo)

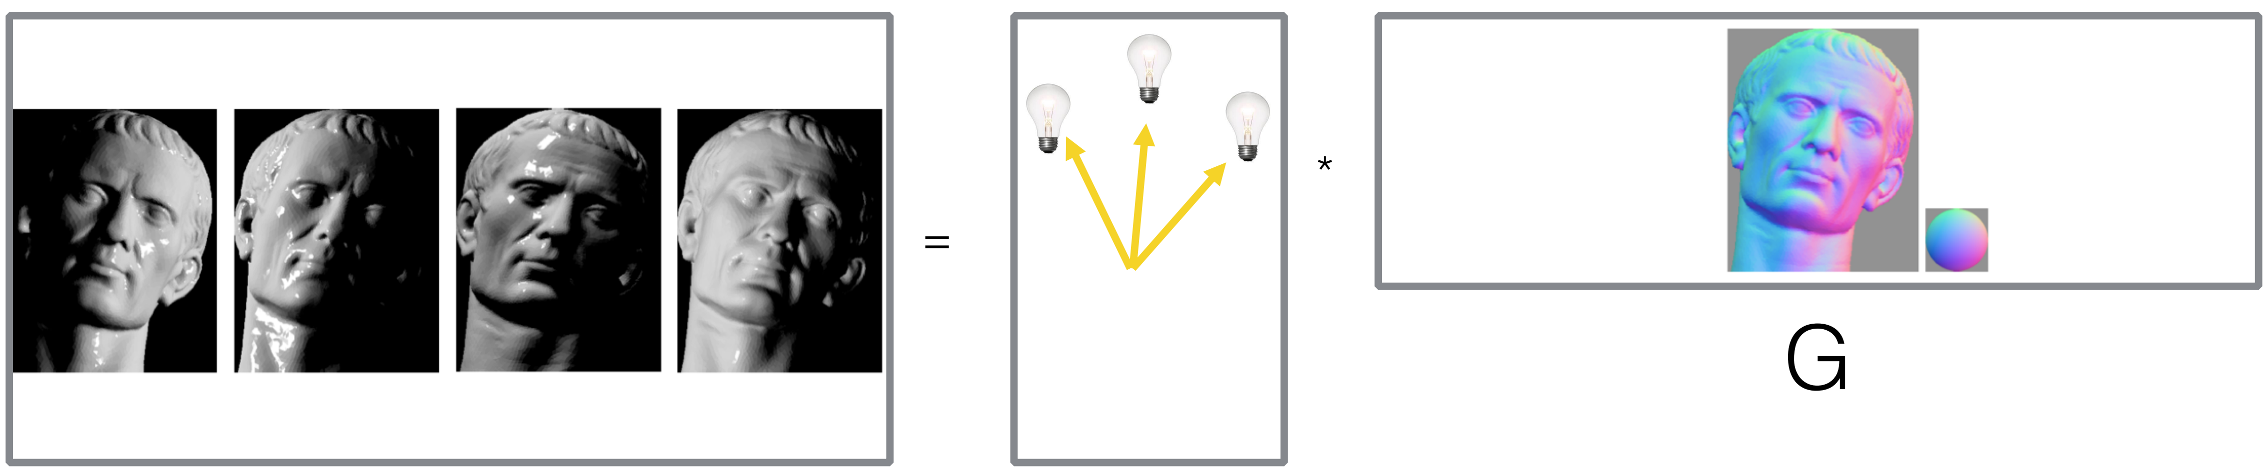

G

I

L^T

Both L and G are now unknown!

This is a matrix factorization problem.

Unknown Lighting

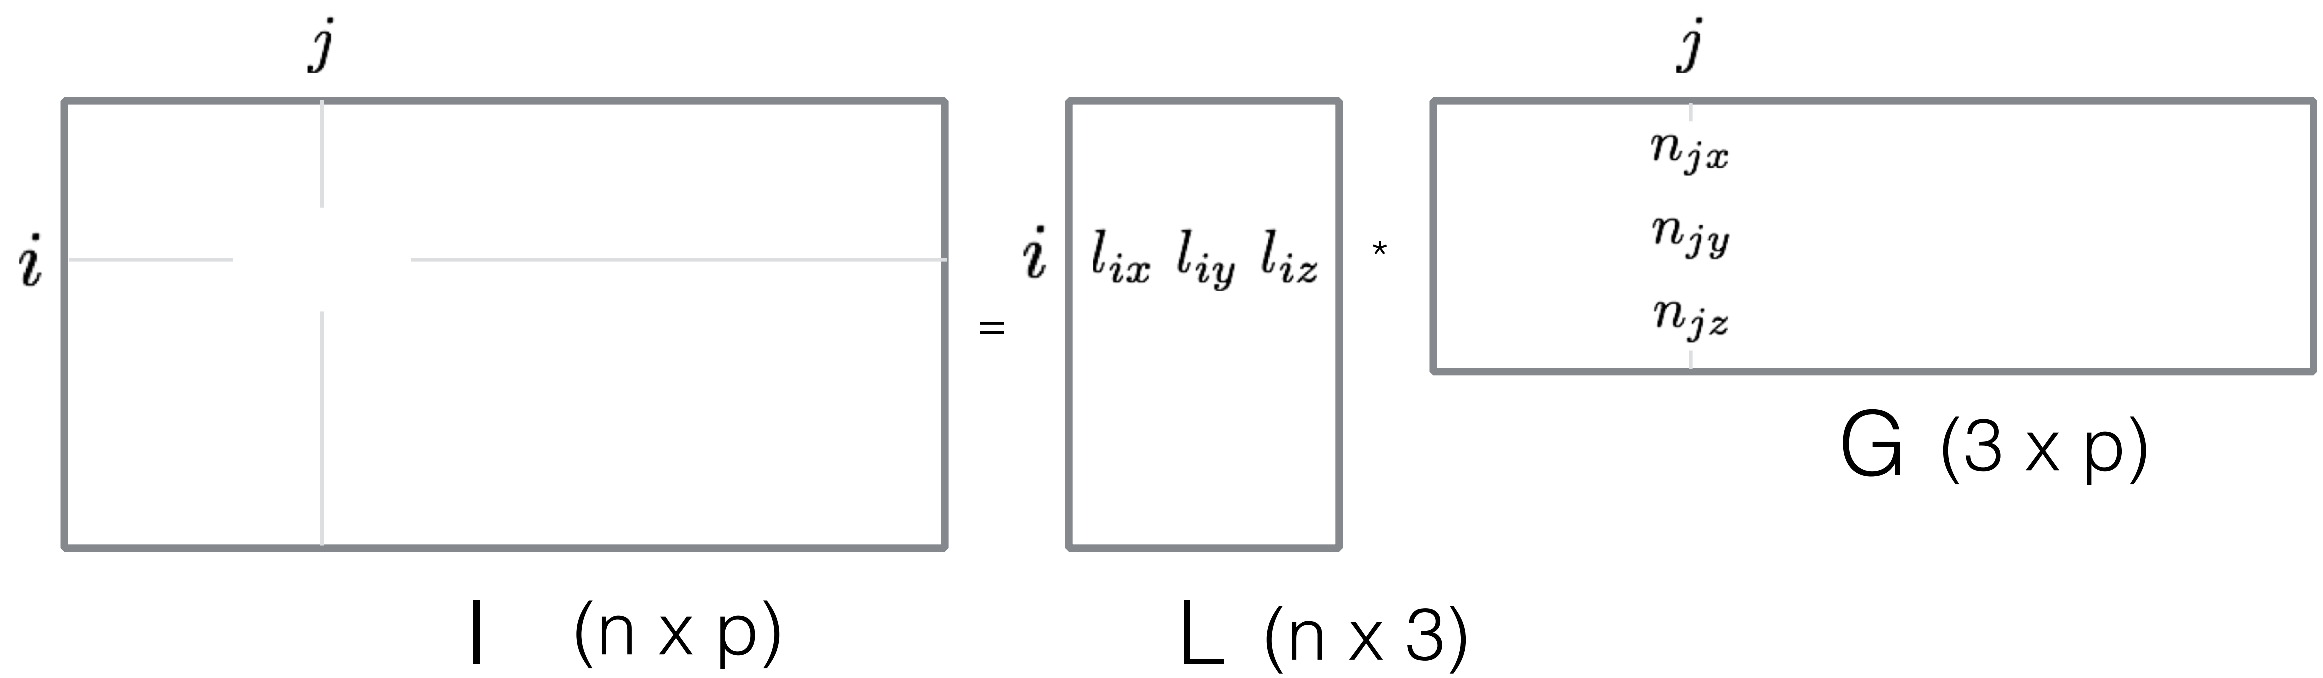

There's hope: We know that I is rank 3

Unknown Lighting

Use the SVD to decompose I:

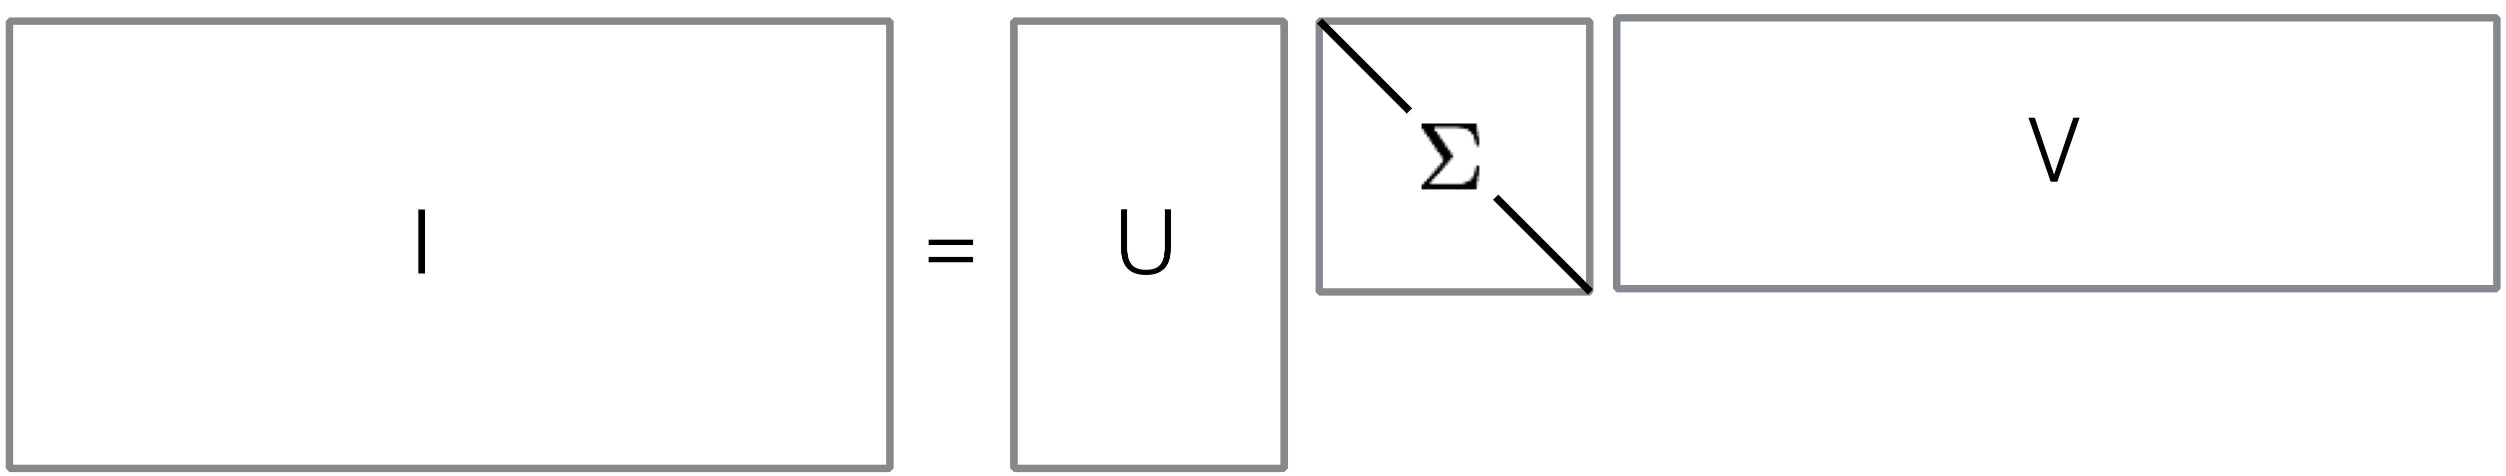

SVD gives the best rank-3 approximation of a matrix.

Unknown Lighting

Use the SVD to decompose I:

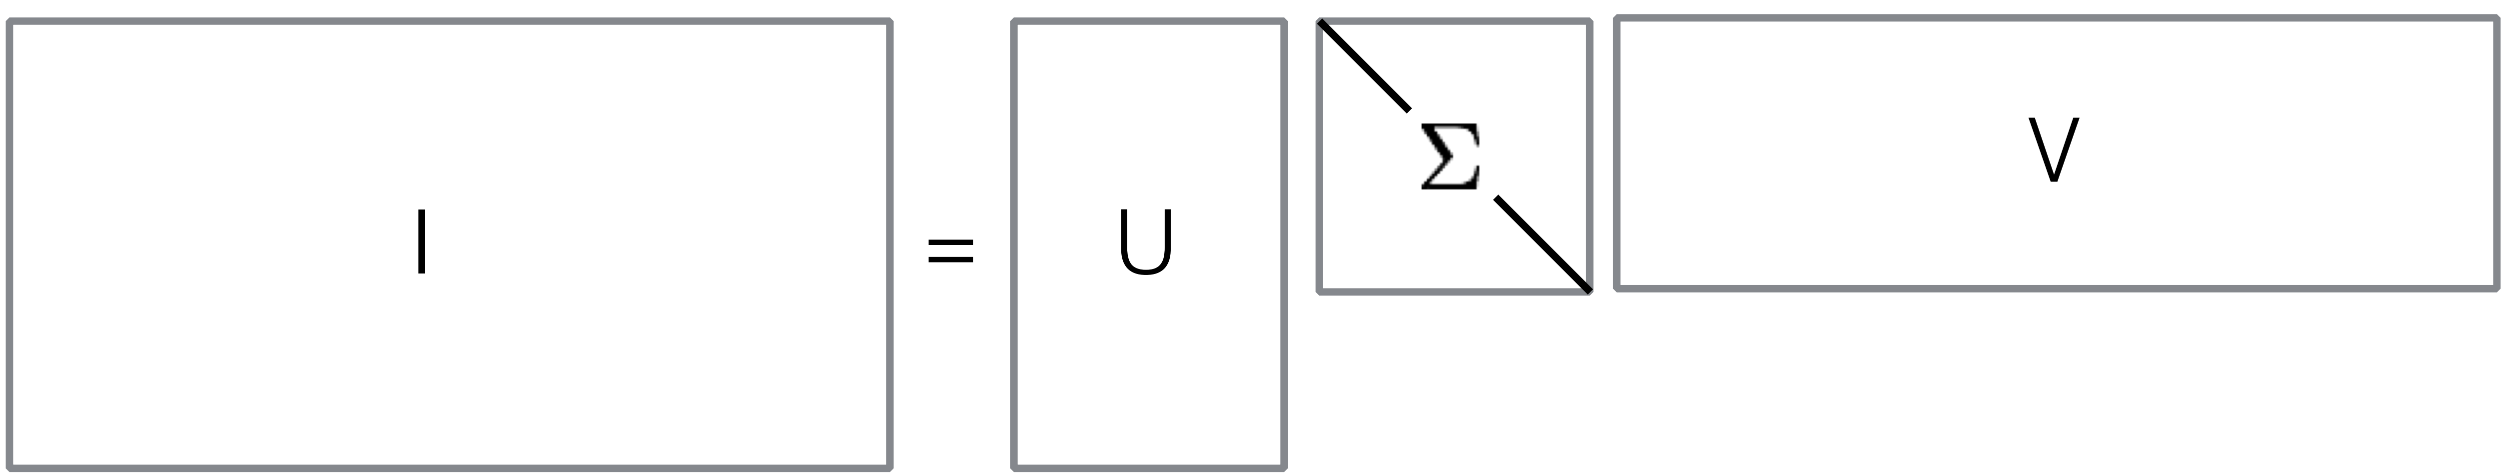

SVD gives the best rank-3 approximation of a matrix.

What do we do with Σ?

Unknown Lighting

Use the SVD to decompose I:

$$I = U\sqrt{\Sigma} \quad \sqrt{\Sigma}V$$

Can we just do that?

Unknown Lighting

Use the SVD to decompose I:

$$I = U\sqrt{\Sigma} \; A \; A^{-1} \; \sqrt{\Sigma}V$$

Can we just do that? …almost.

The decomposition is unique up to an invertible 3x3 A.

Unknown Lighting

Use the SVD to decompose I:

$$I = U\sqrt{\Sigma} \quad A \quad A^{-1} \quad \sqrt{\Sigma}V$$

Can we just do that? …almost. $L = U\sqrt{\Sigma}A, G = A^{-1}\sqrt{\Sigma}V$

The decomposition is unique up to an invertible 3x3 A.

Unknown Lighting

Use the SVD to decompose I:

$$I = U\sqrt{\Sigma} \; A \; A^{-1} \; \sqrt{\Sigma}V$$

$$L = U\sqrt{\Sigma}A, G = A^{-1}\sqrt{\Sigma}V$$

You can find A if you know

- 6 points with the same reflectance, or
- 6 lights with the same intensity.

Unknown Lighting: Ambiguities

- Multiple combinations of lighting and geometry can produce the same sets of images.
- Add assumptions or prior knowledge about geometry or lighting, etc. to limit the ambiguity.

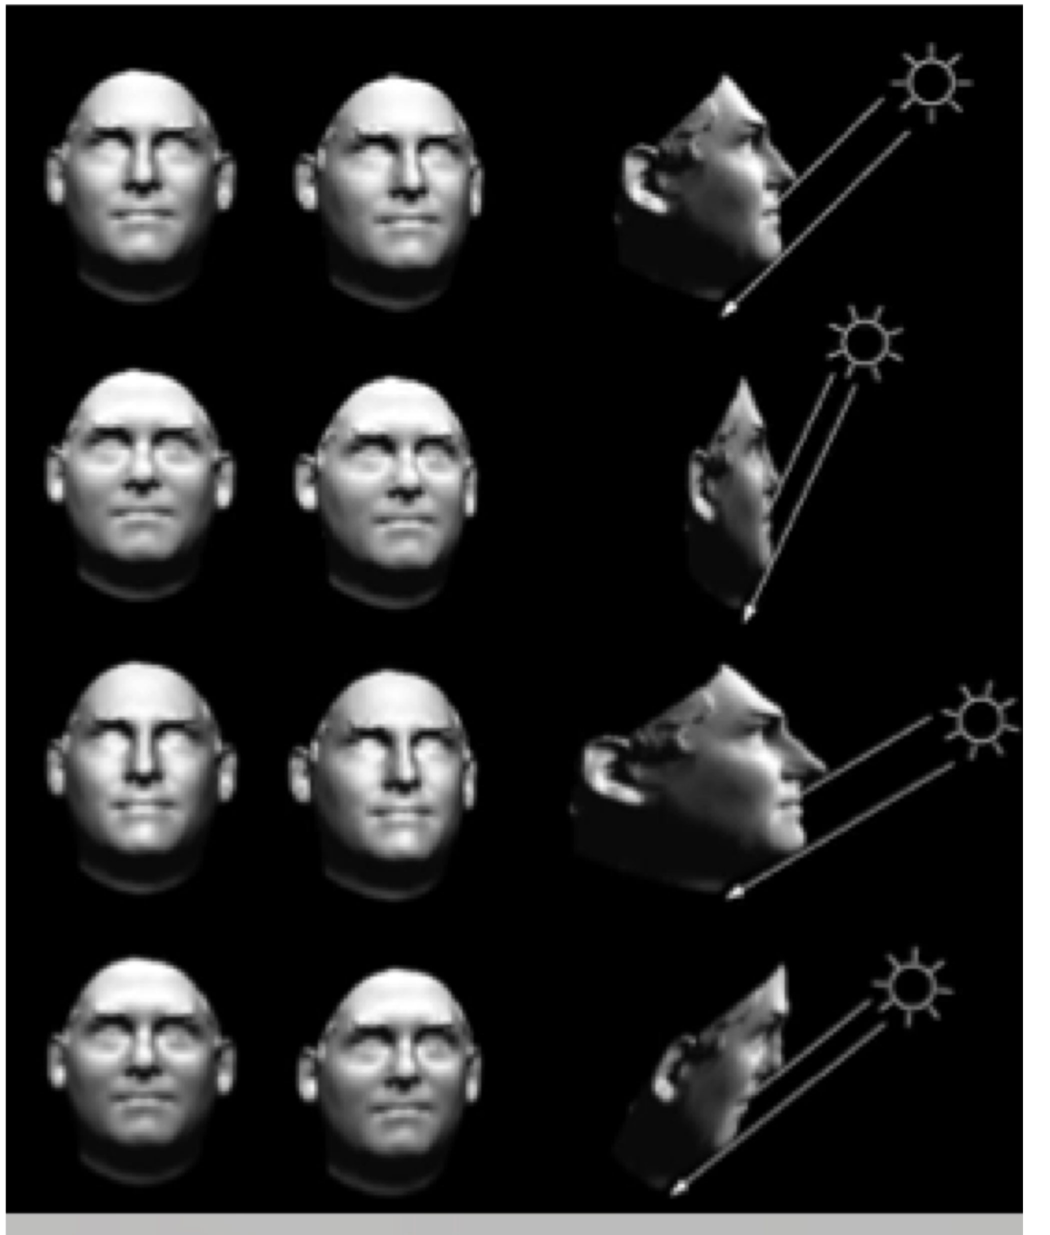

[Belhumeur et al.'97]

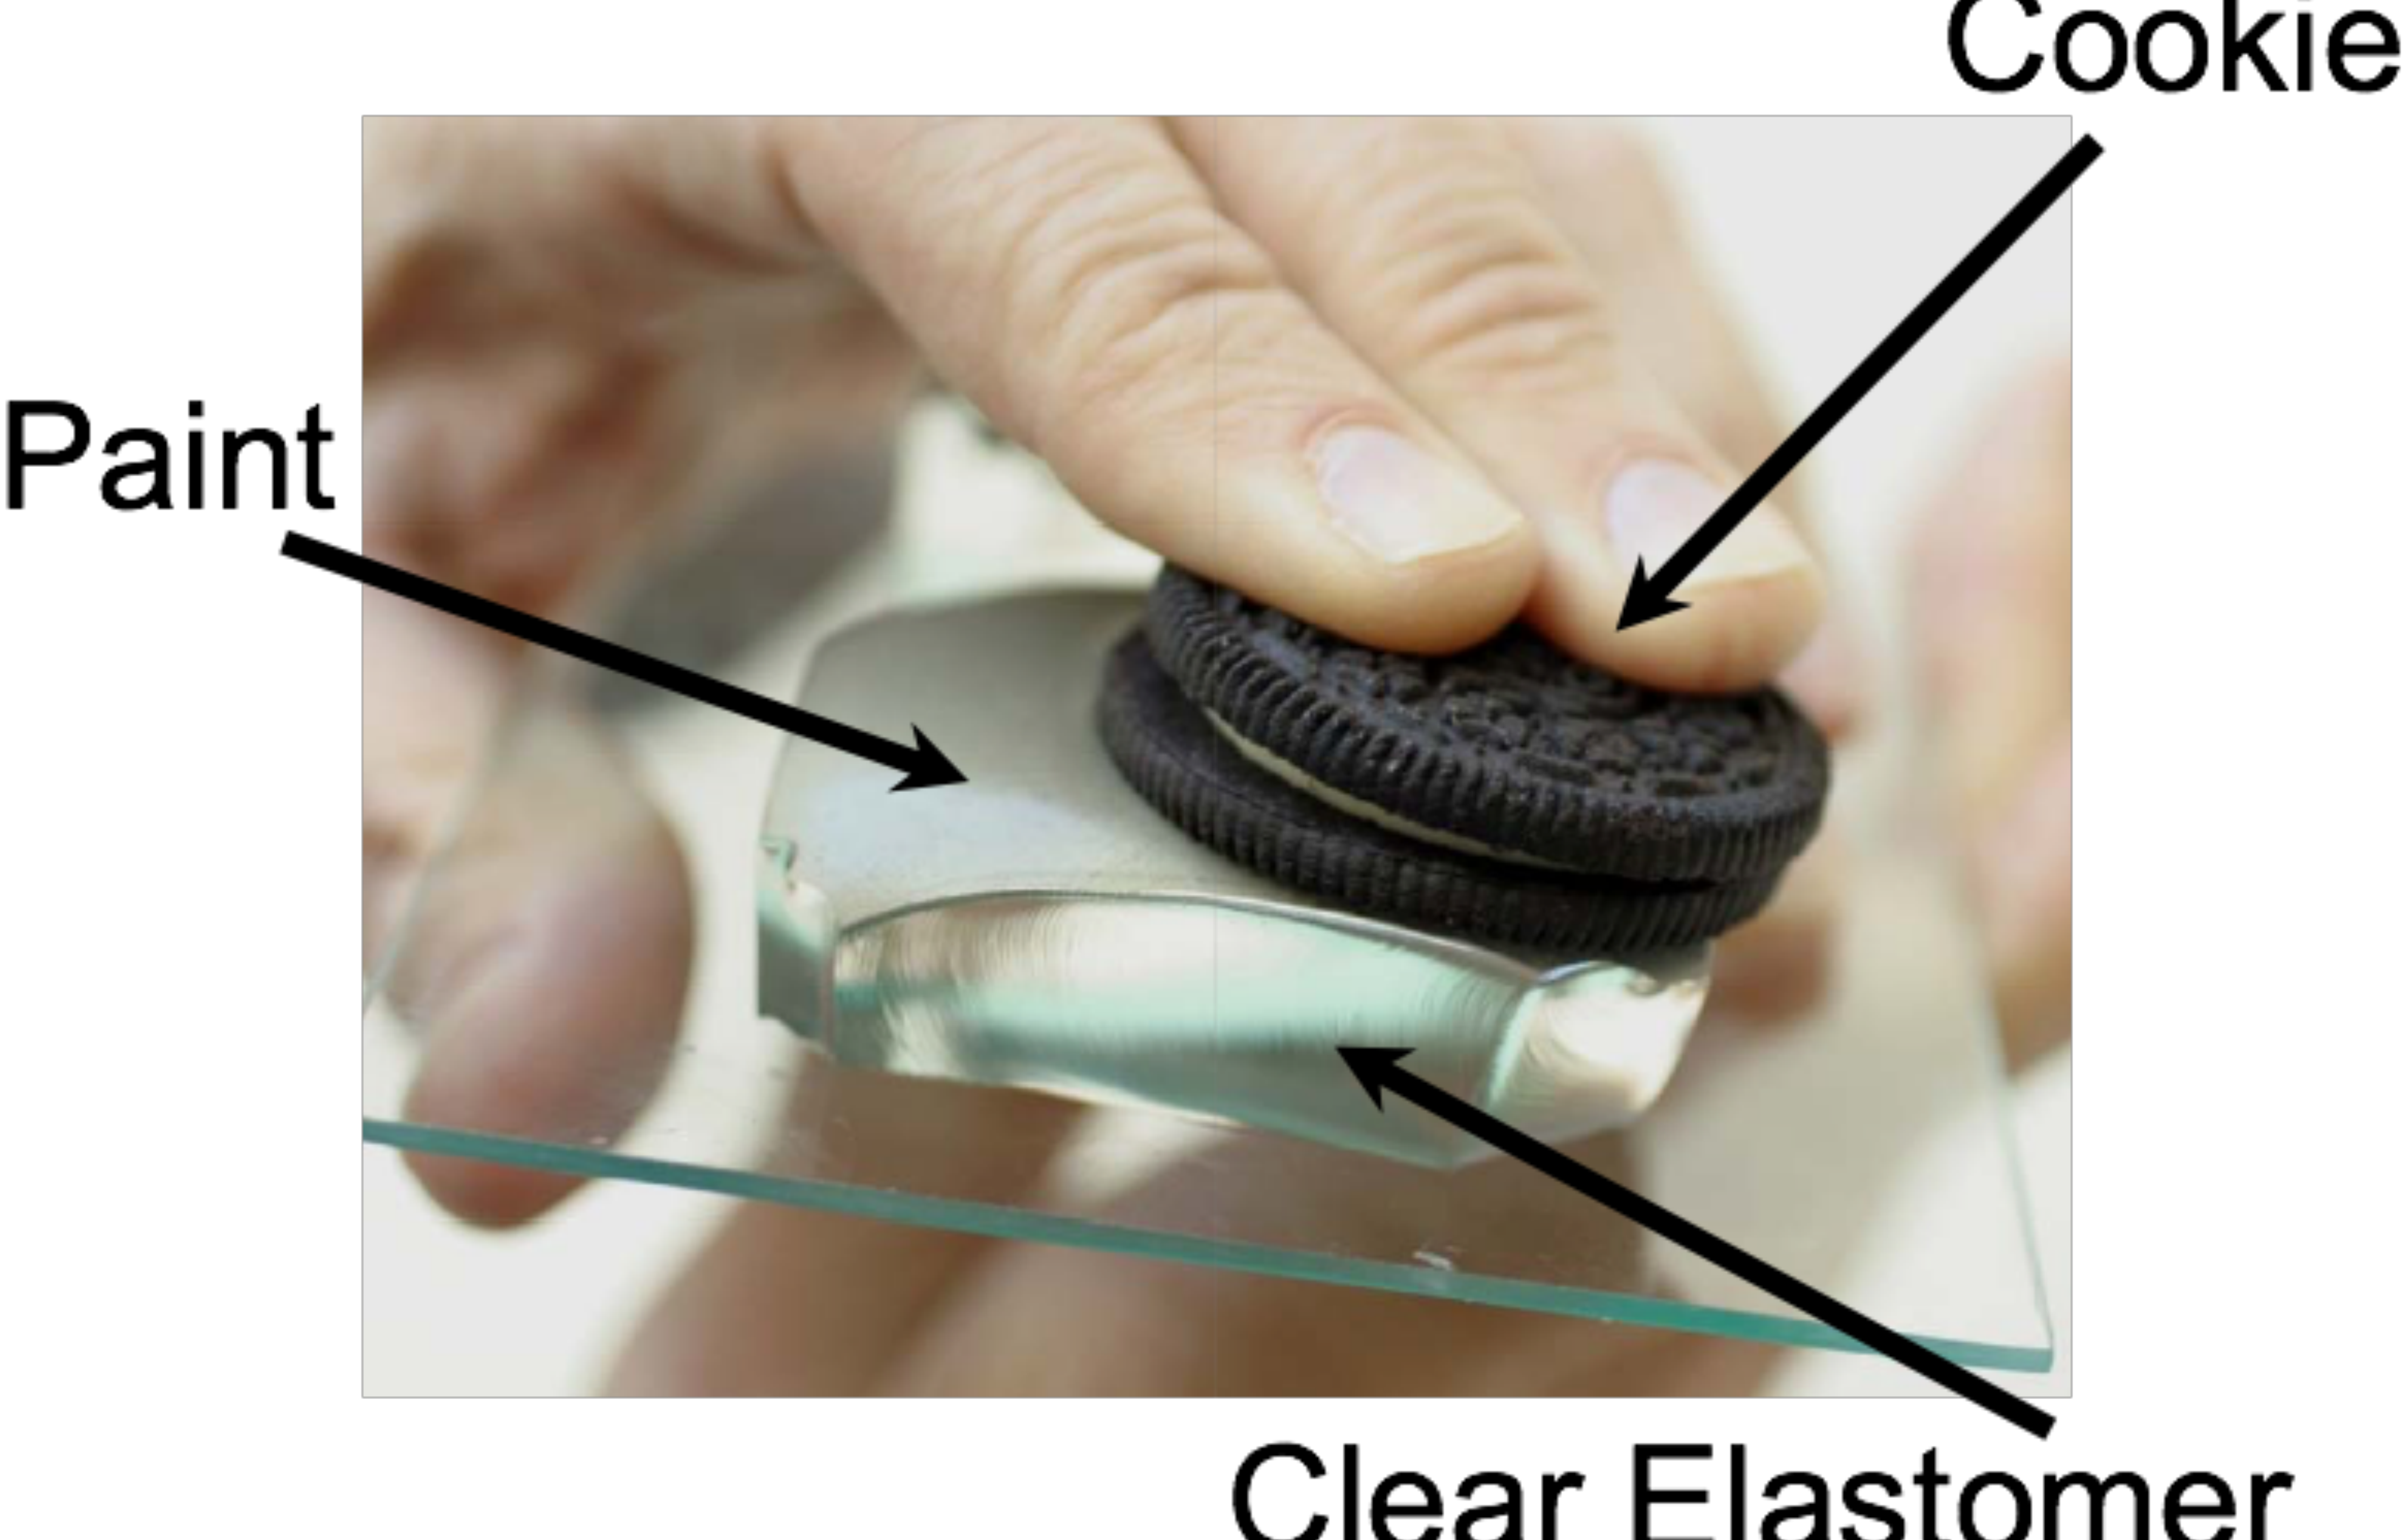

Johnson and Adelson, 2009


OREO

Lights, camera, action

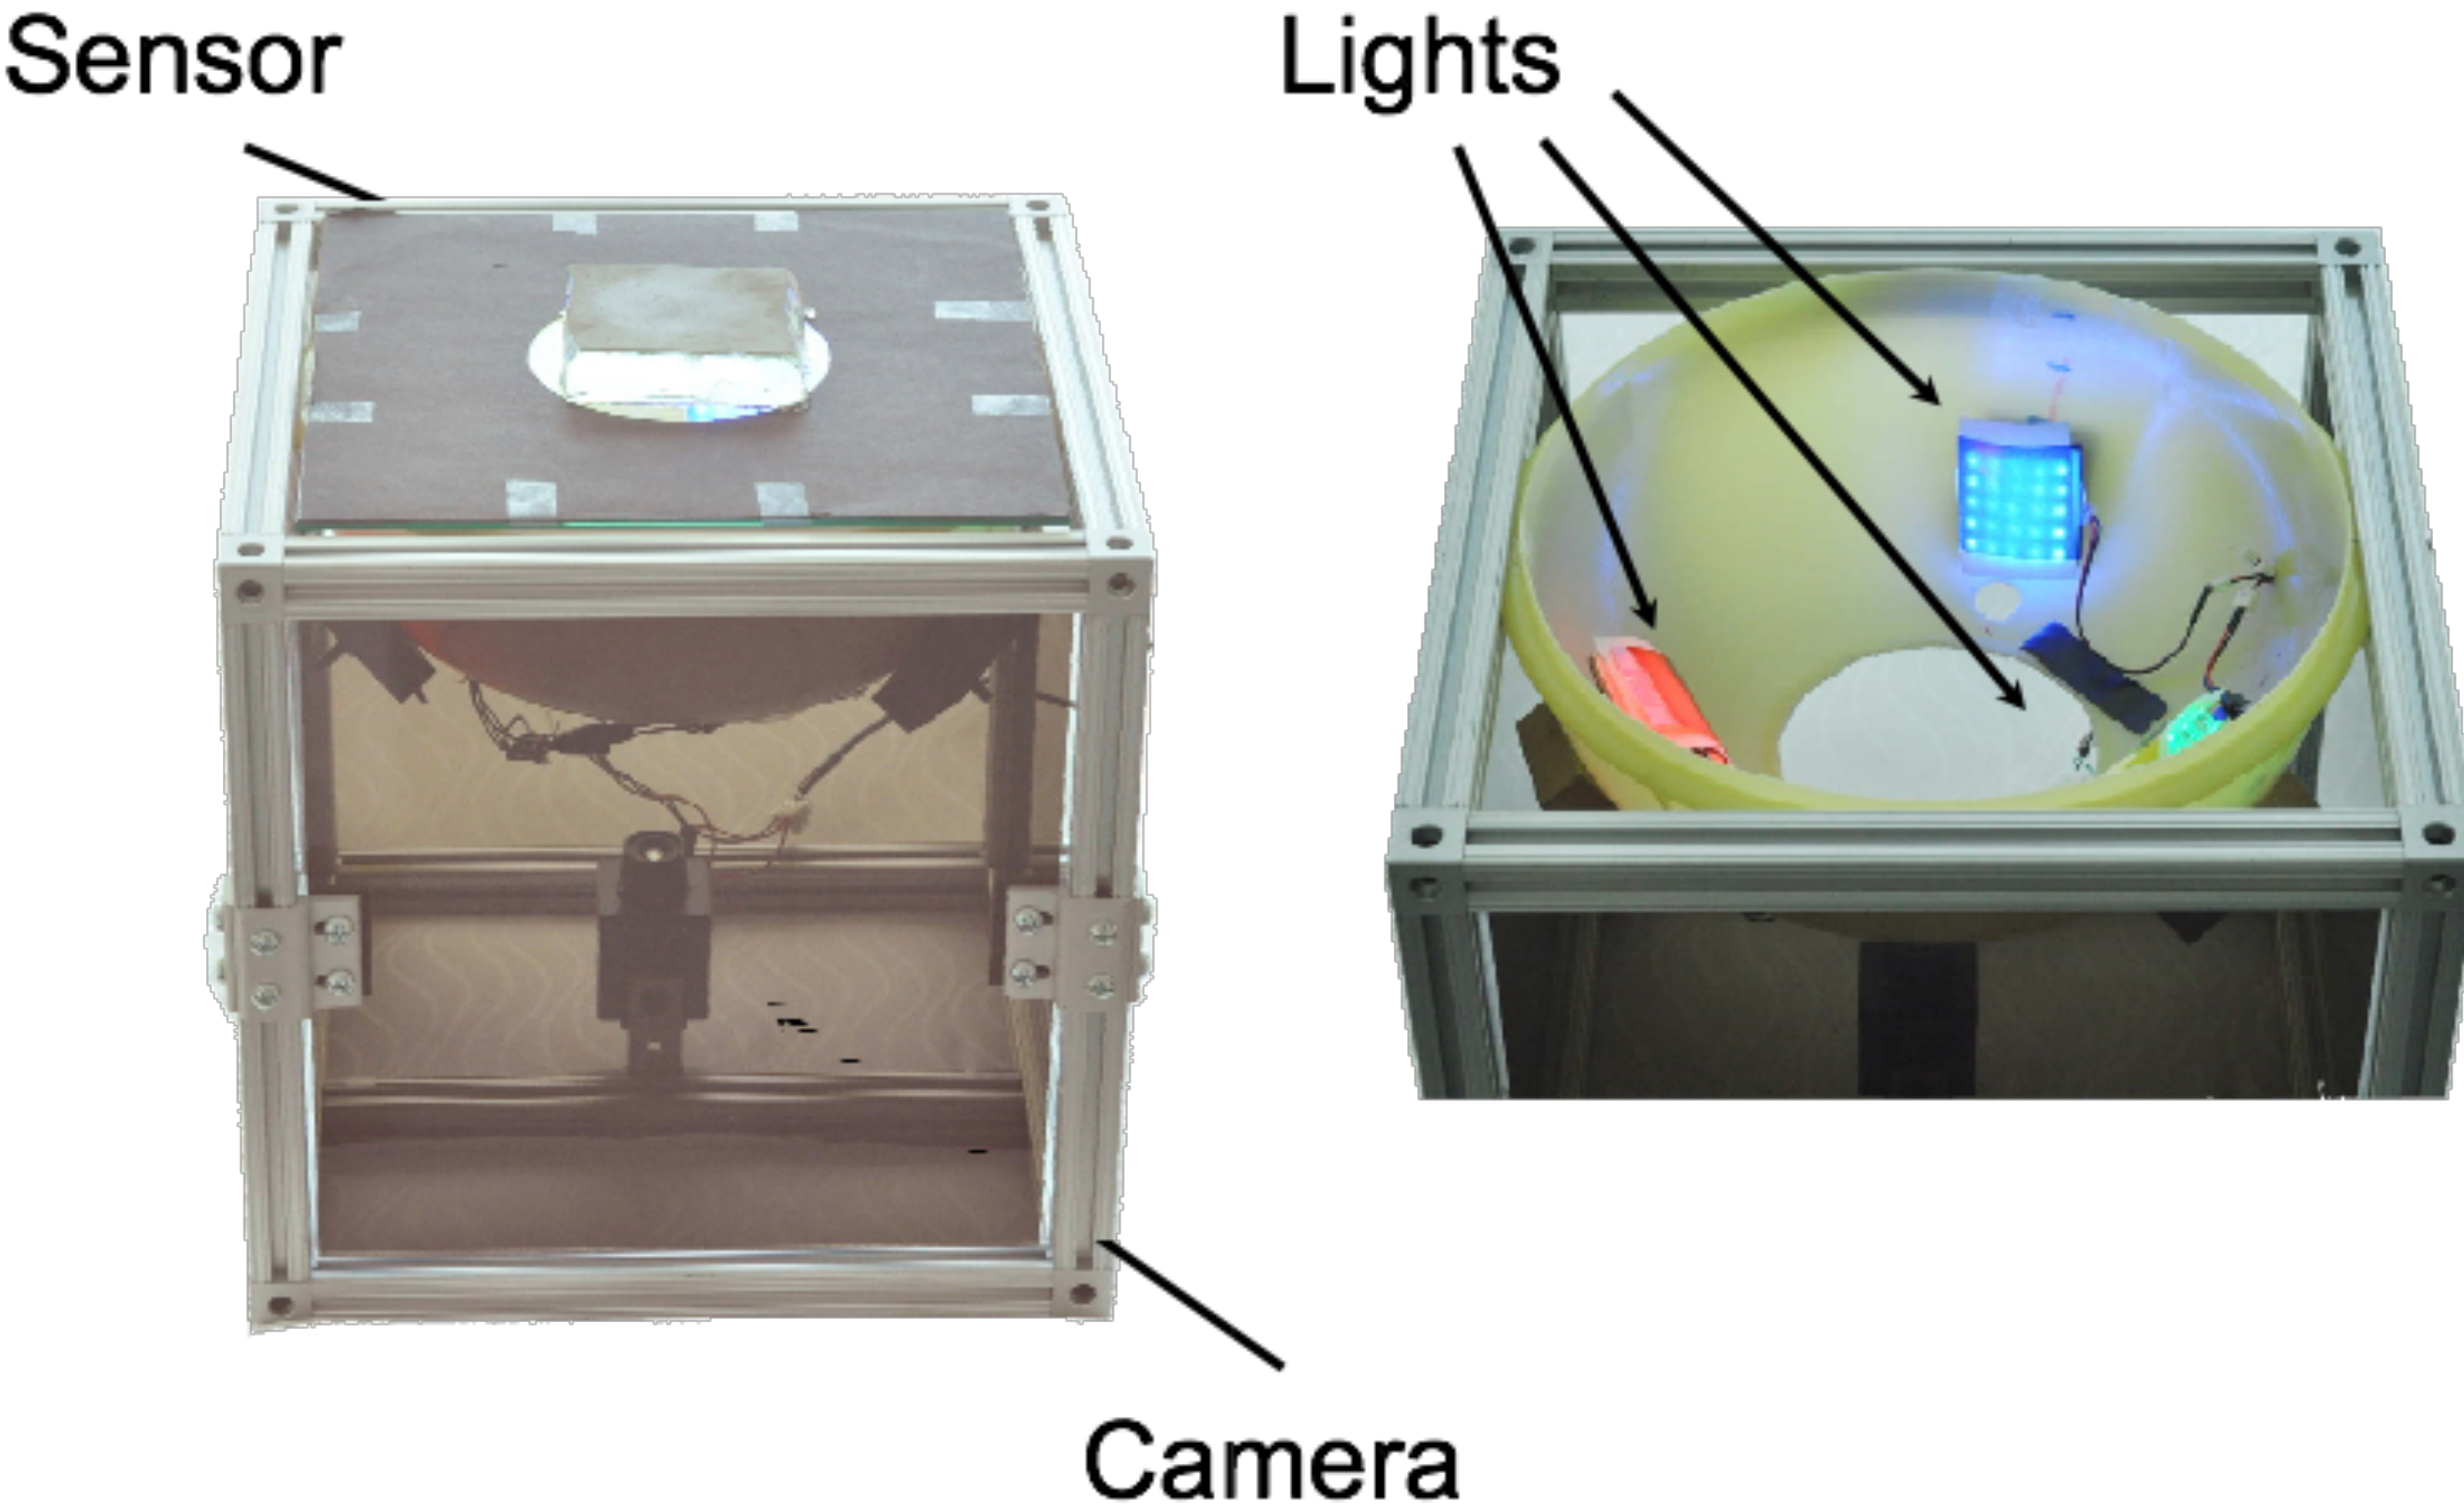

(a) (b)

Figure 2. (a) This decorative pin consists of a glass bas-relief portrait mounted in a shiny gold setting. (b) The RGB image provided by the retrographic sensor. The pin is pressed into the elastomer skin, and colored lights illuminate it from three directions.


OREO

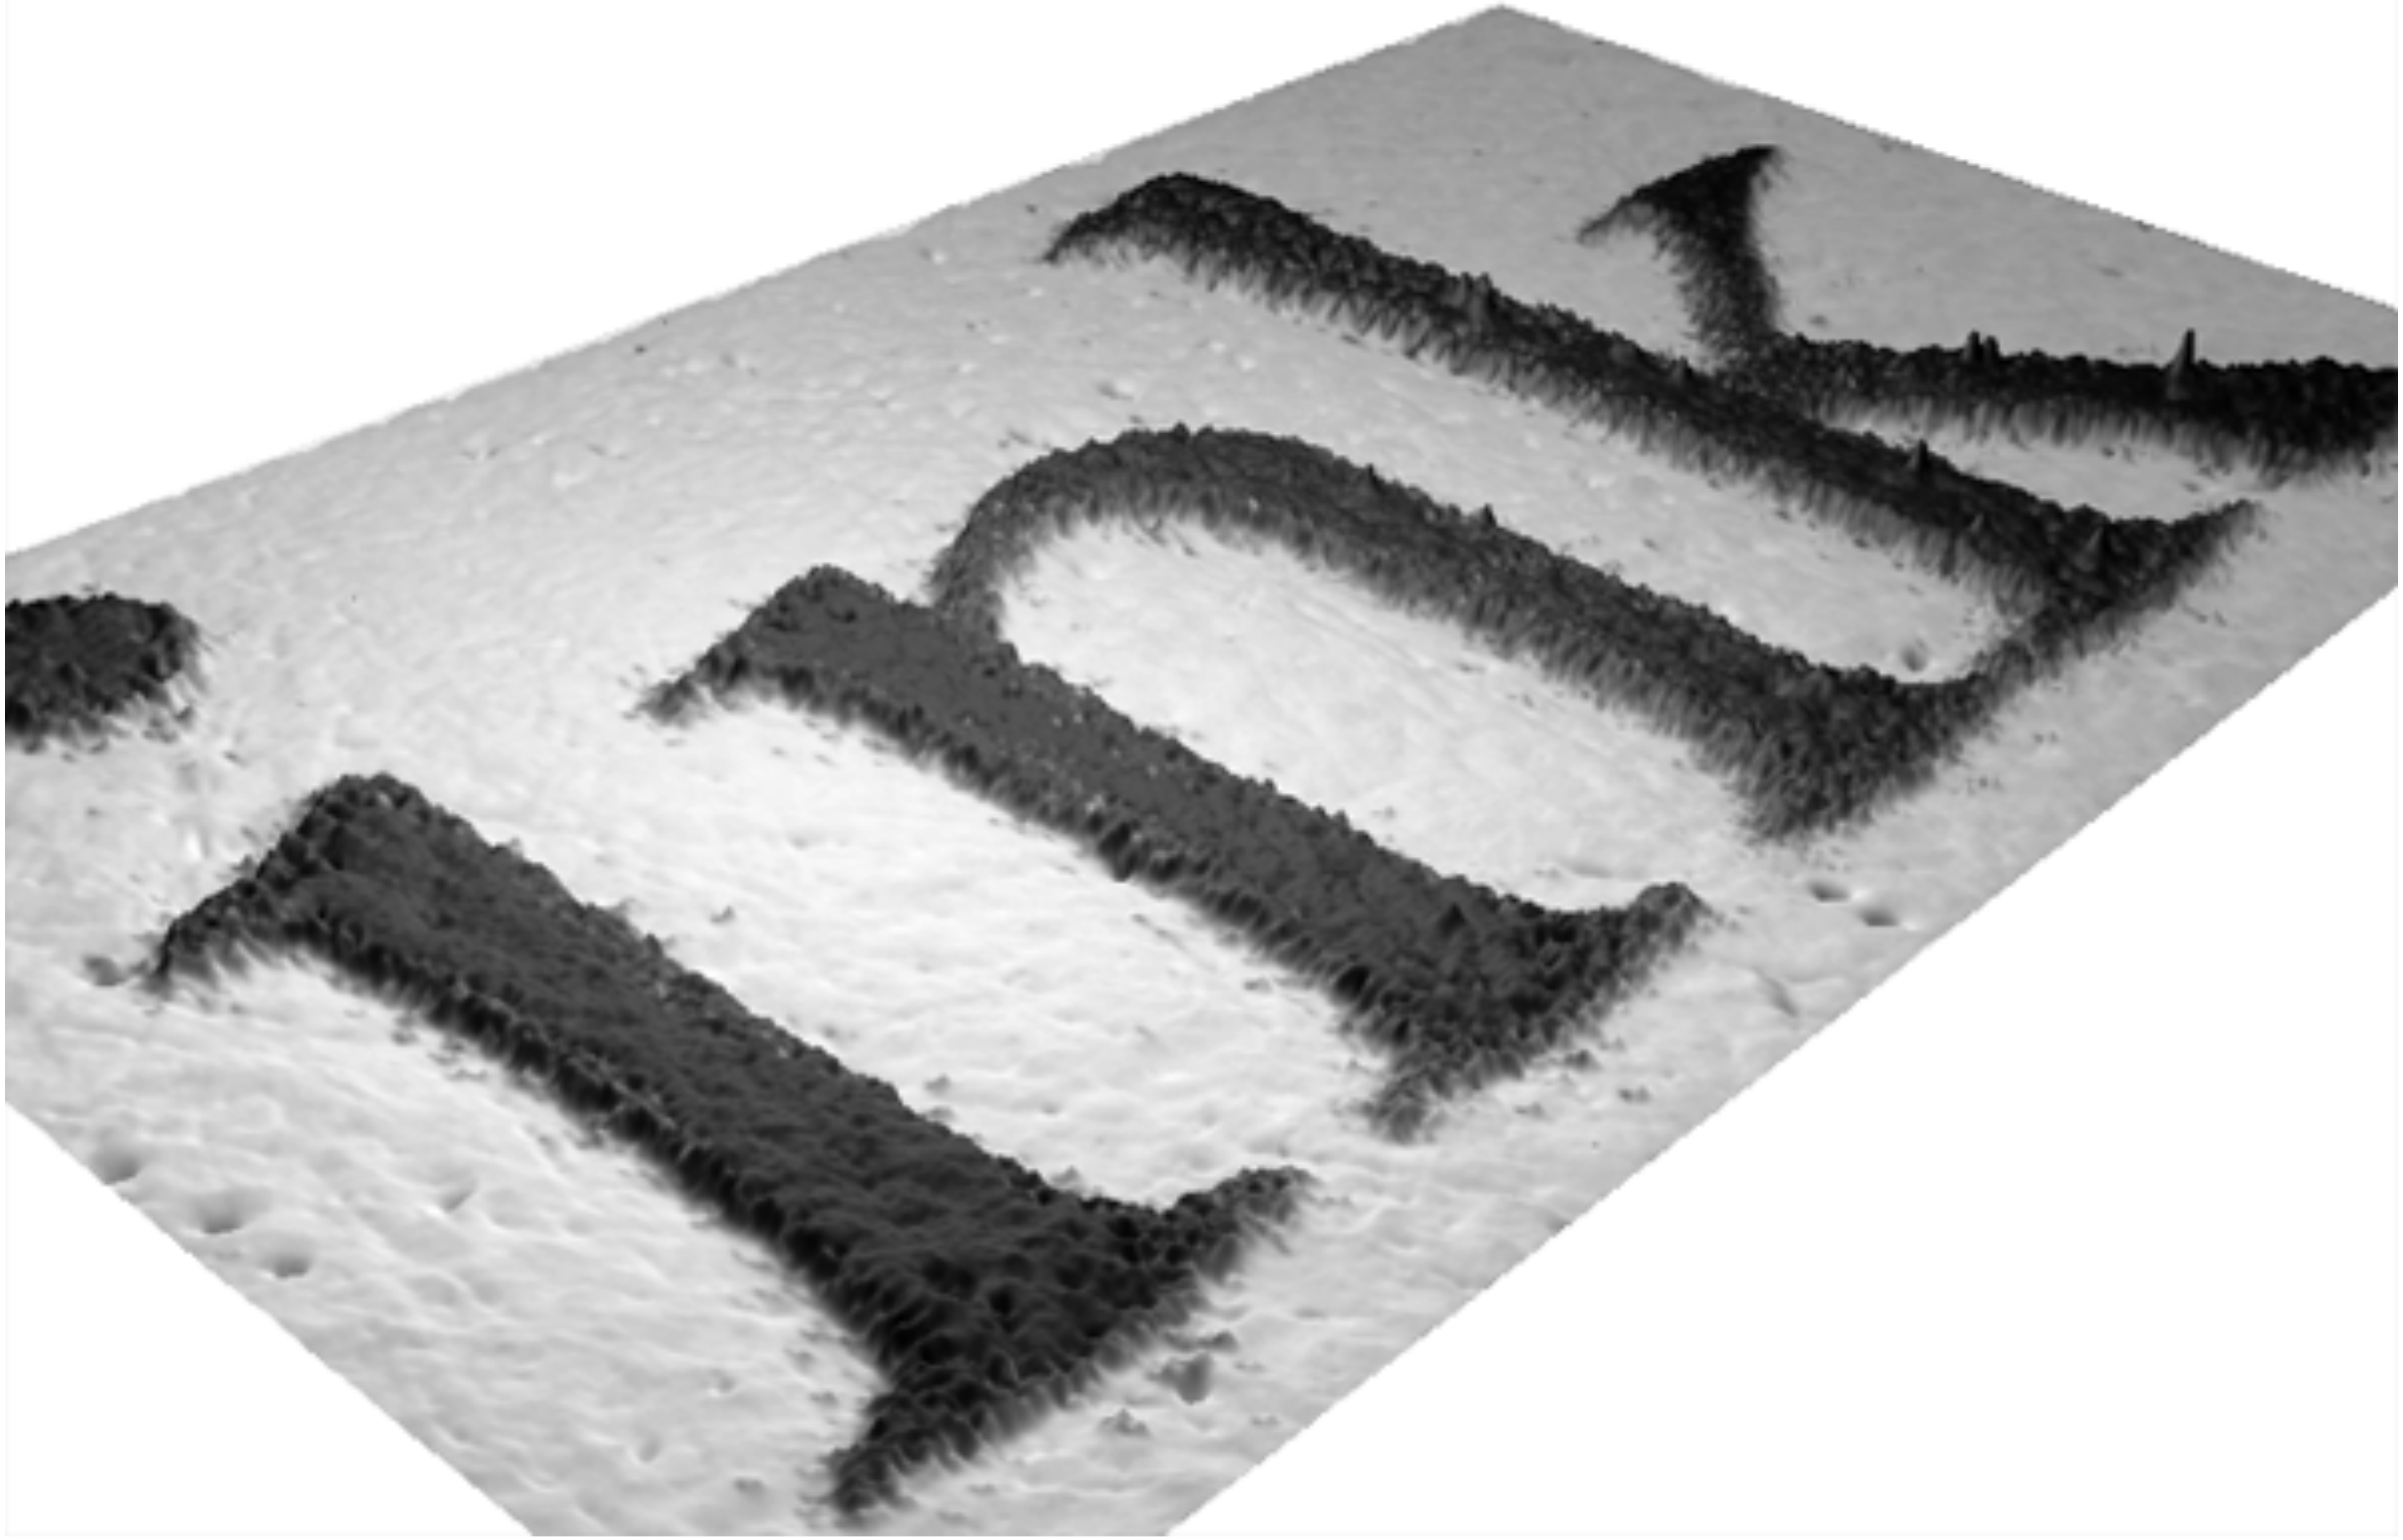

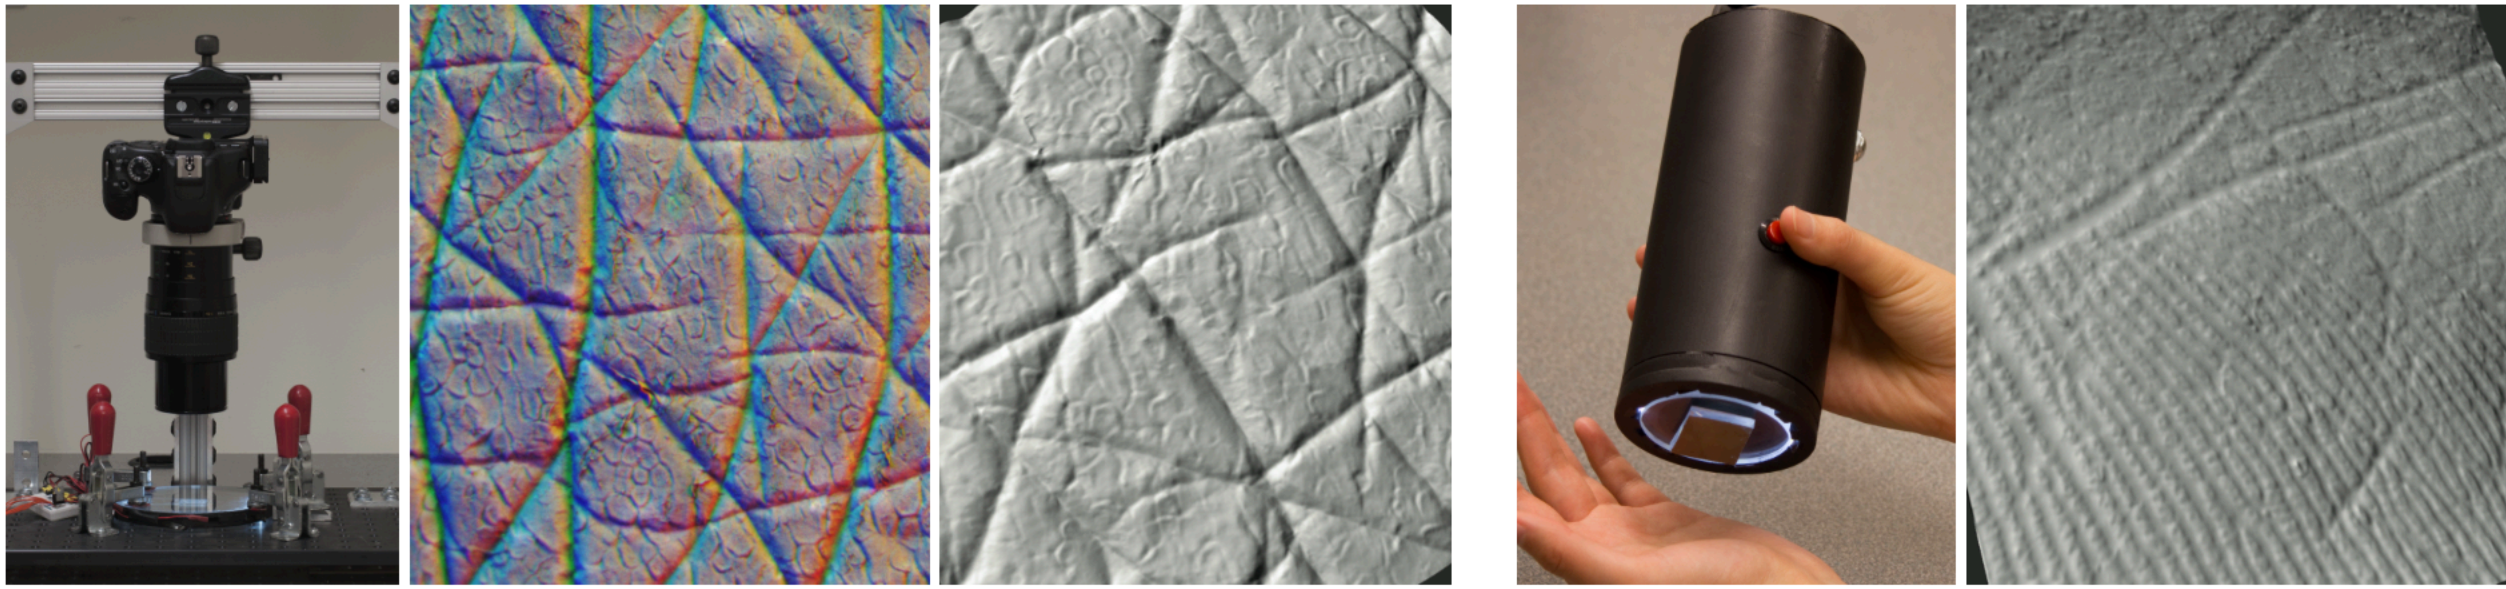

(a) bench configuration (b) captured (c) reconstruction (d) portable configuration (e) reconstruction

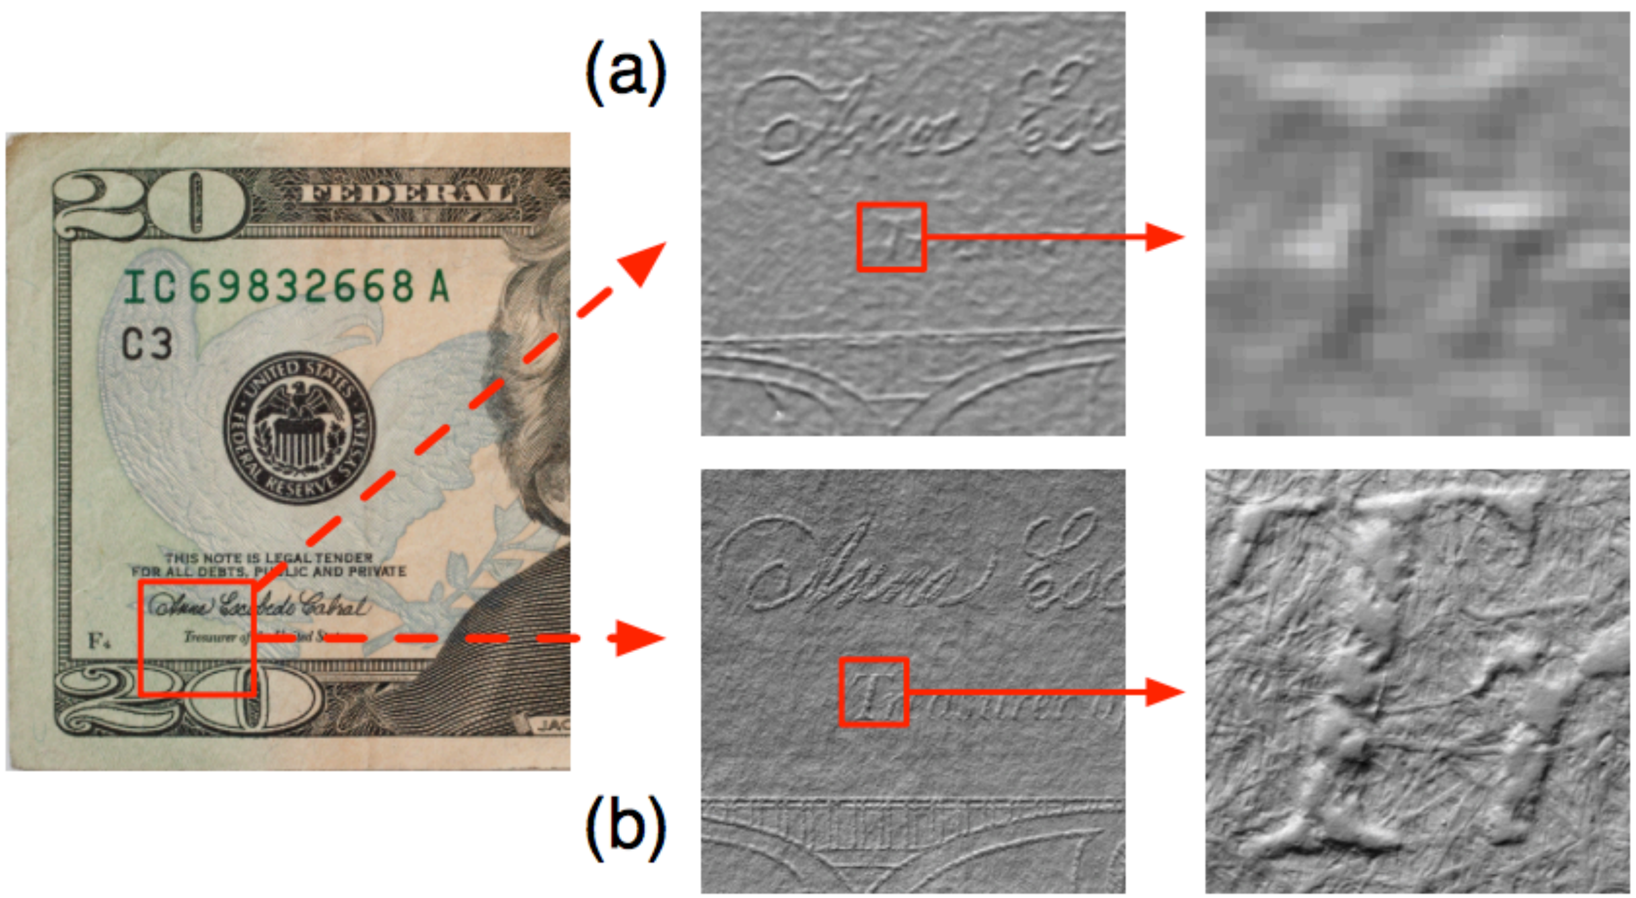

Figure 7: *Comparison with the high-resolution result from the original retrographic sensor. (a) Rendering of the high-resolution \$20 bill example from the original retrographic sensor with a close-up view. (b) Rendering of the captured geometry using our method.*

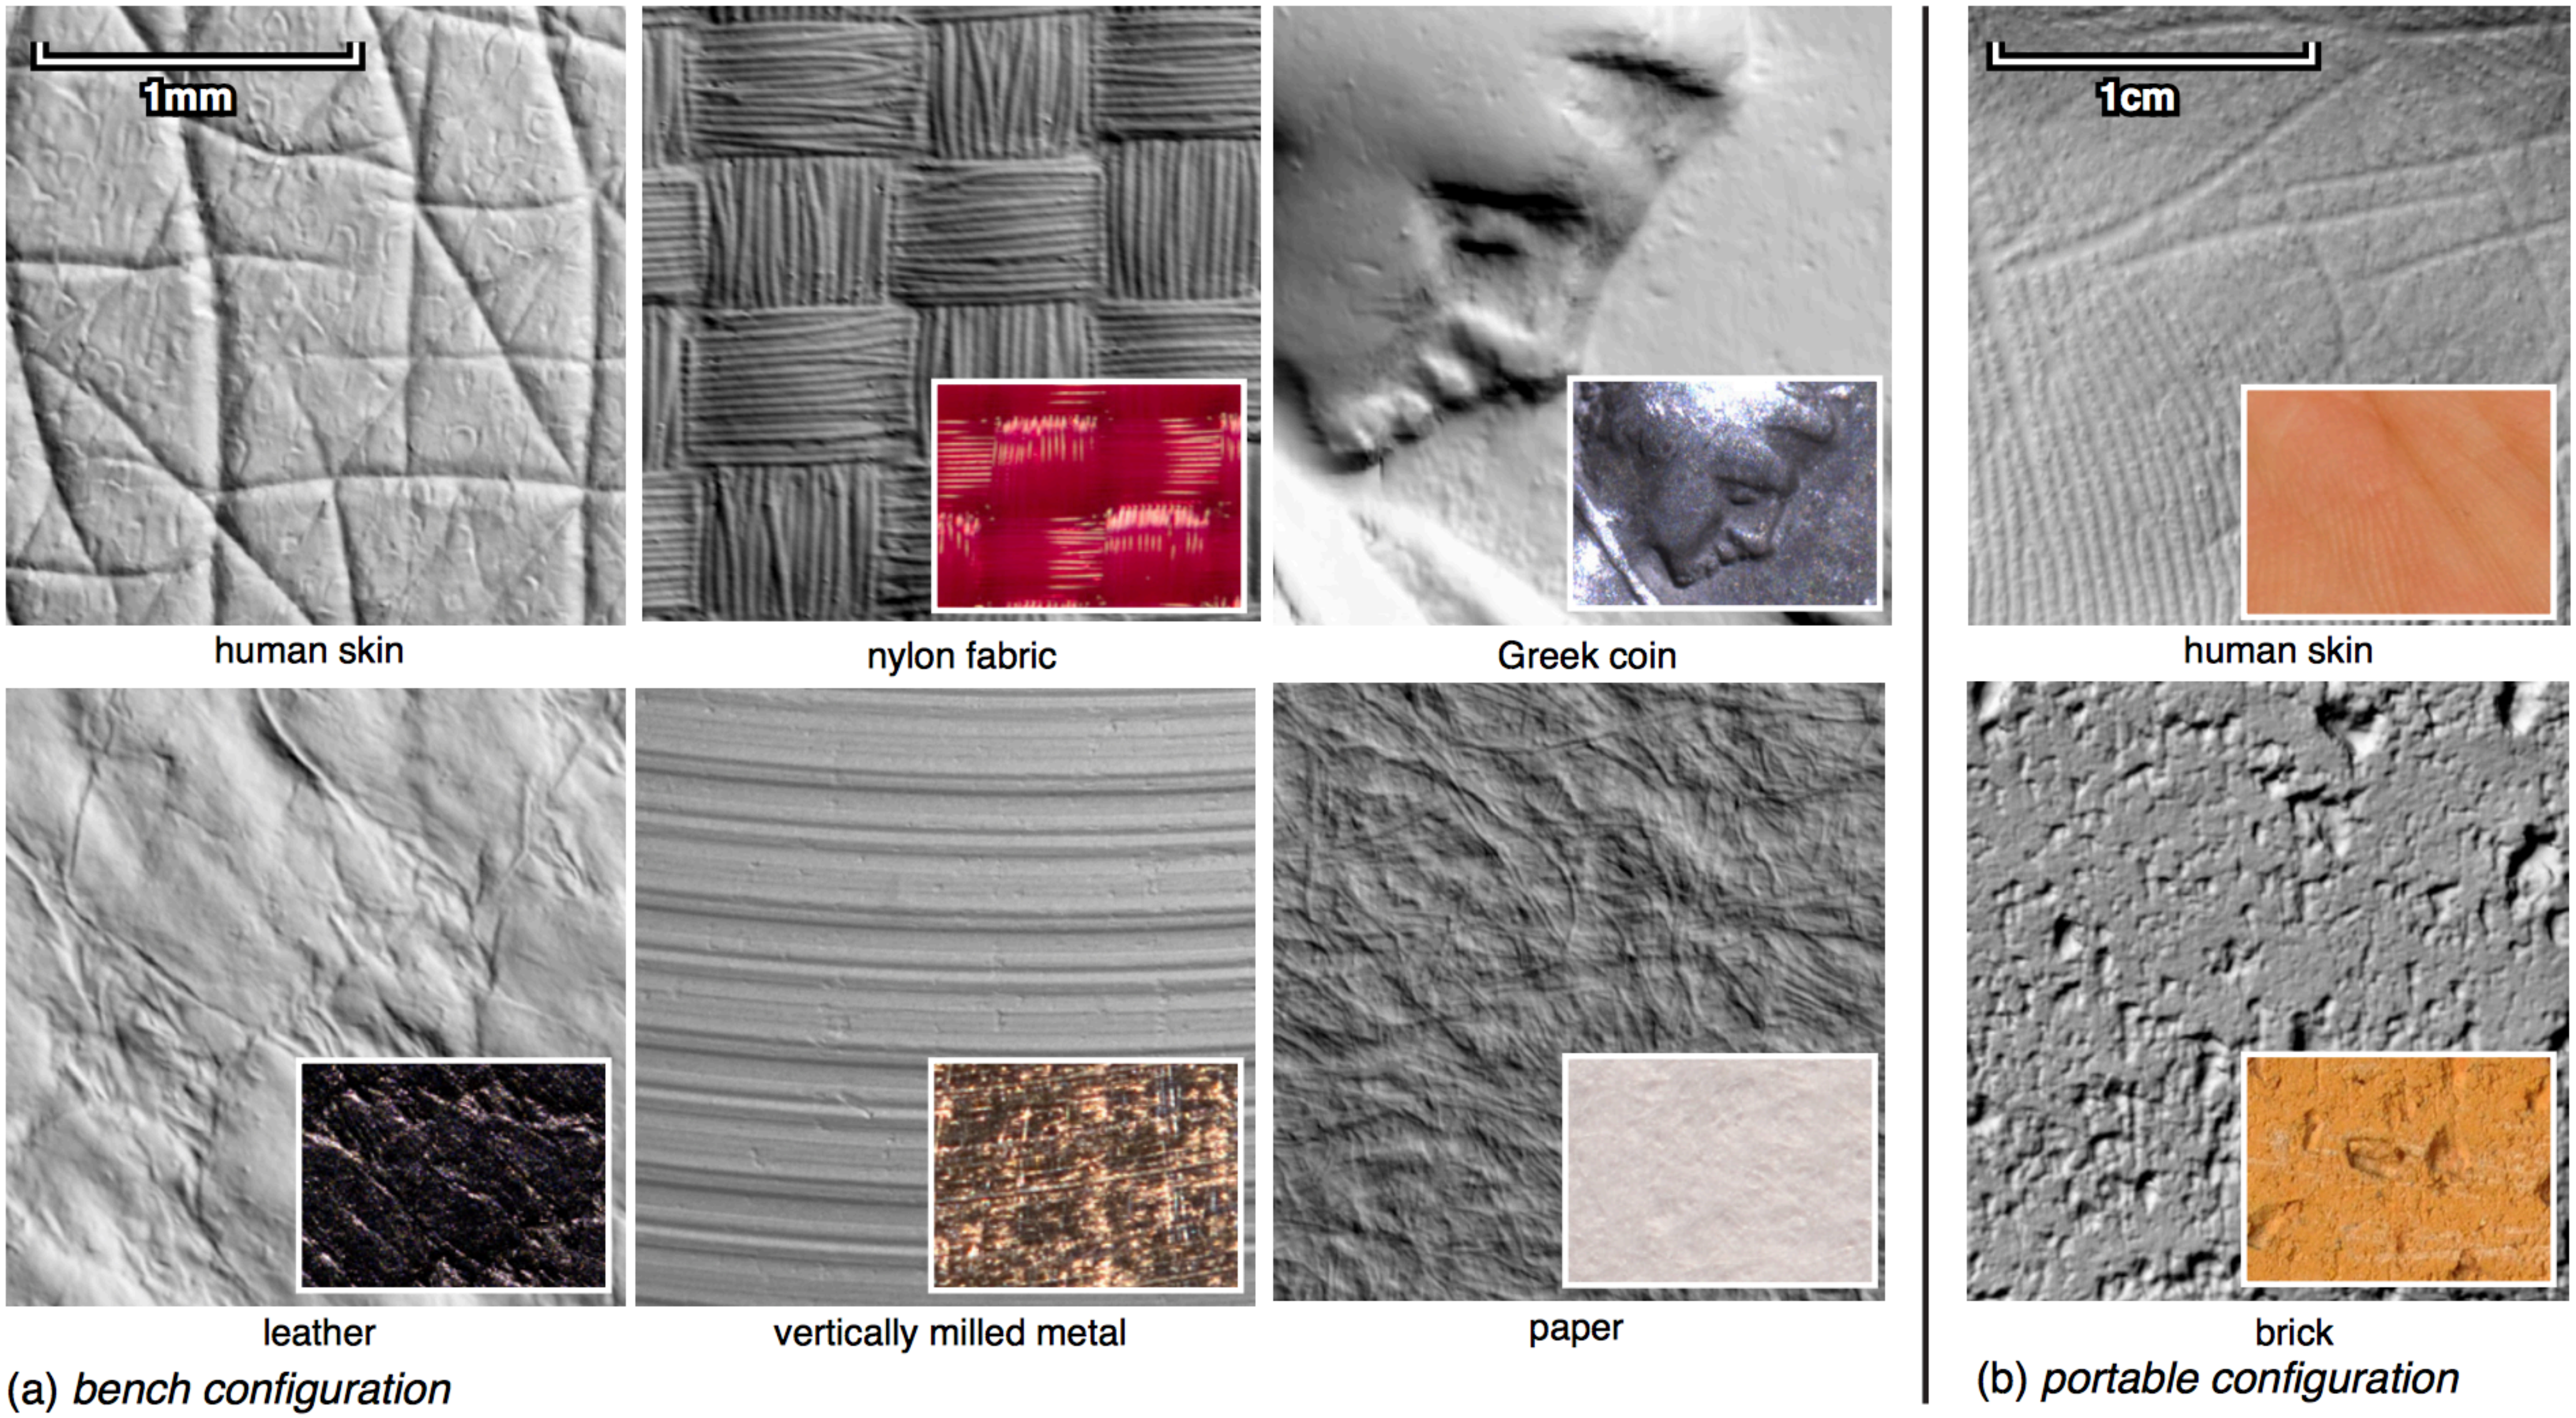

Figure 9: *Example geometry measured with the bench and portable configurations. Outer image: rendering under direct lighting. Inset: macro photograph of original sample. Scale shown in upper left. Color images are shown for context and are to similar, but not exact scale.*