CS4670/5760: Computer Vision

Kavita Bala

Lecture 29: Photometric Stereo 2

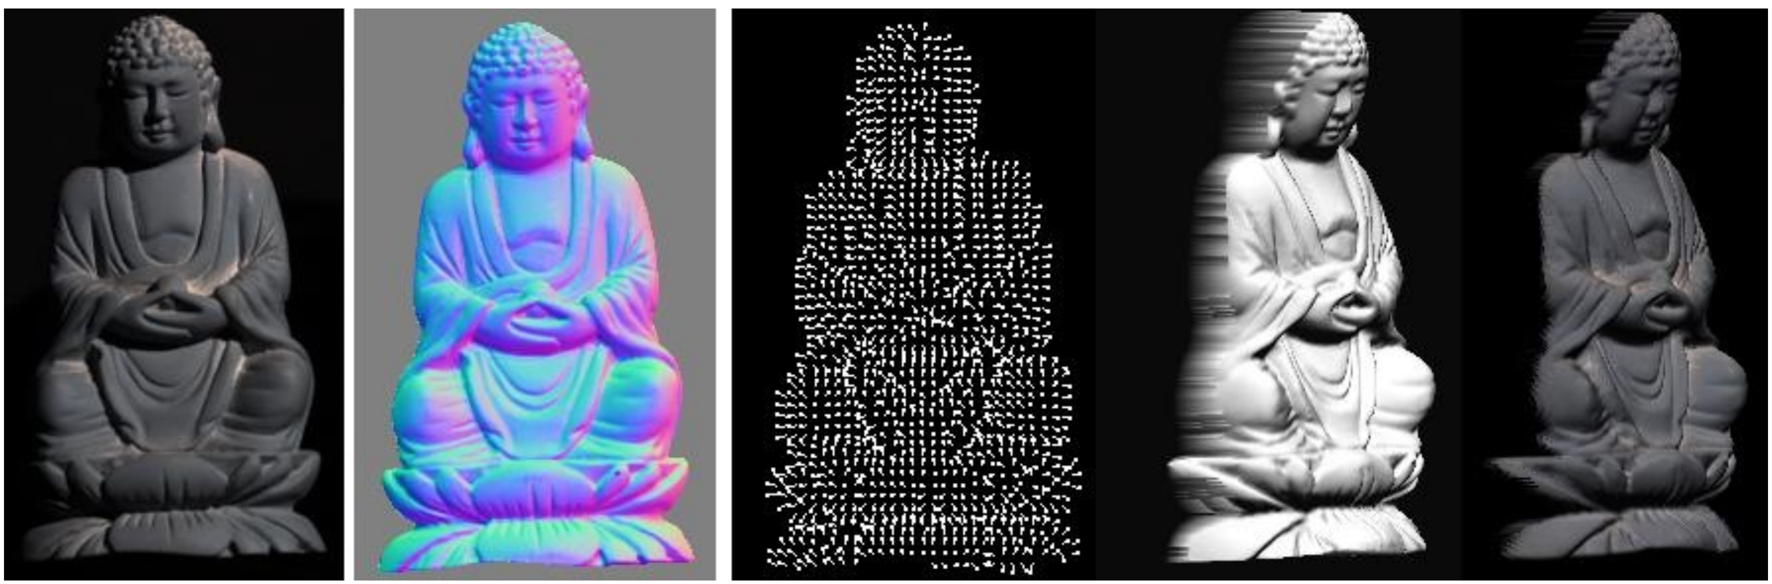

Thanks to Scott Wehrwein

Announcements

- PA 4 out tonight
- HW 2 out tonight
- Wed/Fri: MVS, sFM

Lambertian Reflectance

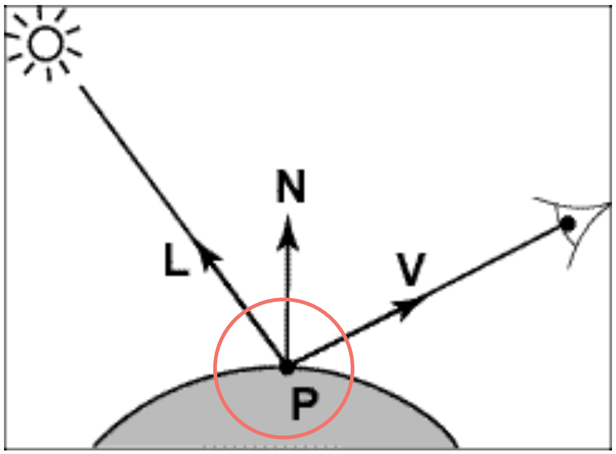

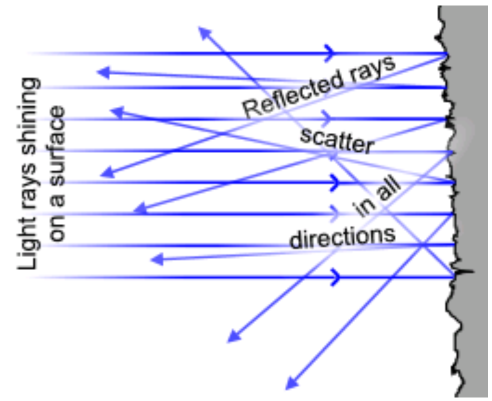

$$I = N \cdot L$$

Image intensity = Surface normal • Light direction

Image intensity ∝ cos(angle between N and L)

Solving the Equations

$$\underbrace{\begin{bmatrix} I_1 & I_2 & I_3 \end{bmatrix}}_{\substack{\mathbf{I} \\ 1 \times 3}} = \underbrace{k_d \mathbf{N}^T}_{\substack{\mathbf{G} \\ 1 \times 3}} \underbrace{\begin{bmatrix} \mathbf{L_1} & \mathbf{L_2} & \mathbf{L_3} \end{bmatrix}}_{\substack{\mathcal{L} \\ 3 \times 3}}$$

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

- When is L nonsingular (invertible)?
  - >= 3 light directions are linearly independent, or:
  - All light direction vectors cannot lie in a plane.
- What if we have more than one pixel?
  - Stack them all into one big system.

More than Three Lights

$$\begin{bmatrix} I_1 & \dots & I_n \end{bmatrix} = k_d \mathbf{N}^T \begin{bmatrix} \mathbf{L_1} & \dots & \mathbf{L_n} \end{bmatrix}$$

- Solve using least squares (normal equations):

$$\begin{aligned} \mathbf{I} &= \mathbf{GL} \\ \mathbf{IL^T} &= \mathbf{GLL^T} \\ \mathbf{G} &= (\mathbf{IL^T})(\mathbf{LL^T})^{-1} \end{aligned}$$

- Equivalently use SVD
- Given G, solve for N and k_d as before.

More than one pixel

Stack all pixels into one system:

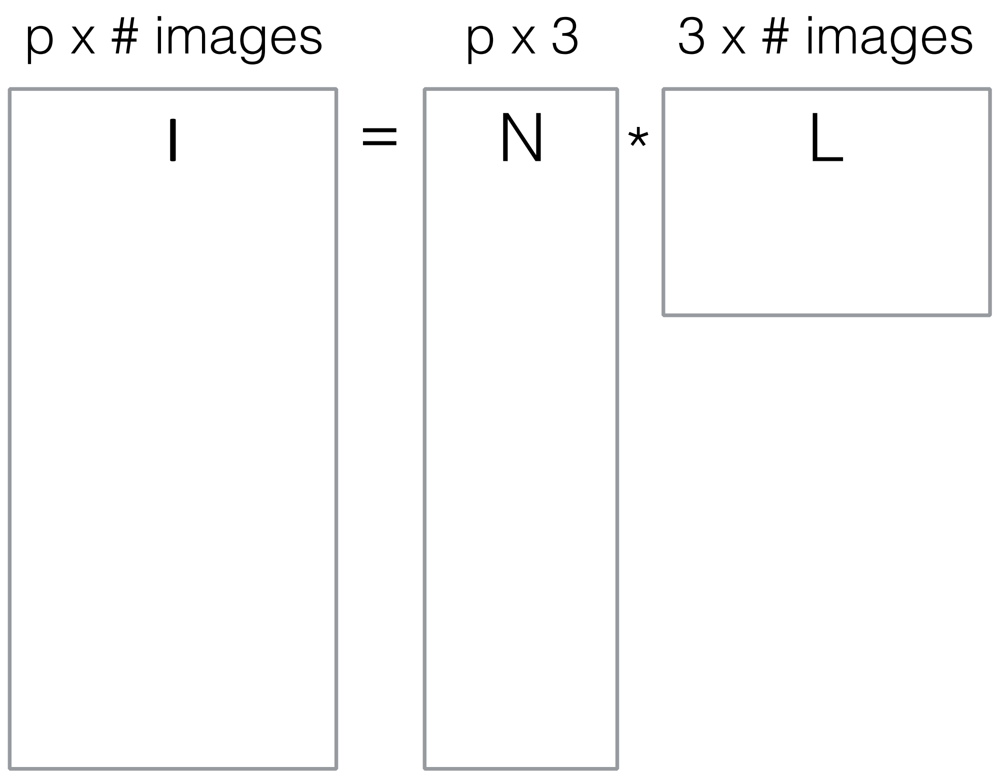

Solve as before.

Depth Map from Normal Map

- We now have a surface normal, but how do we get depth?

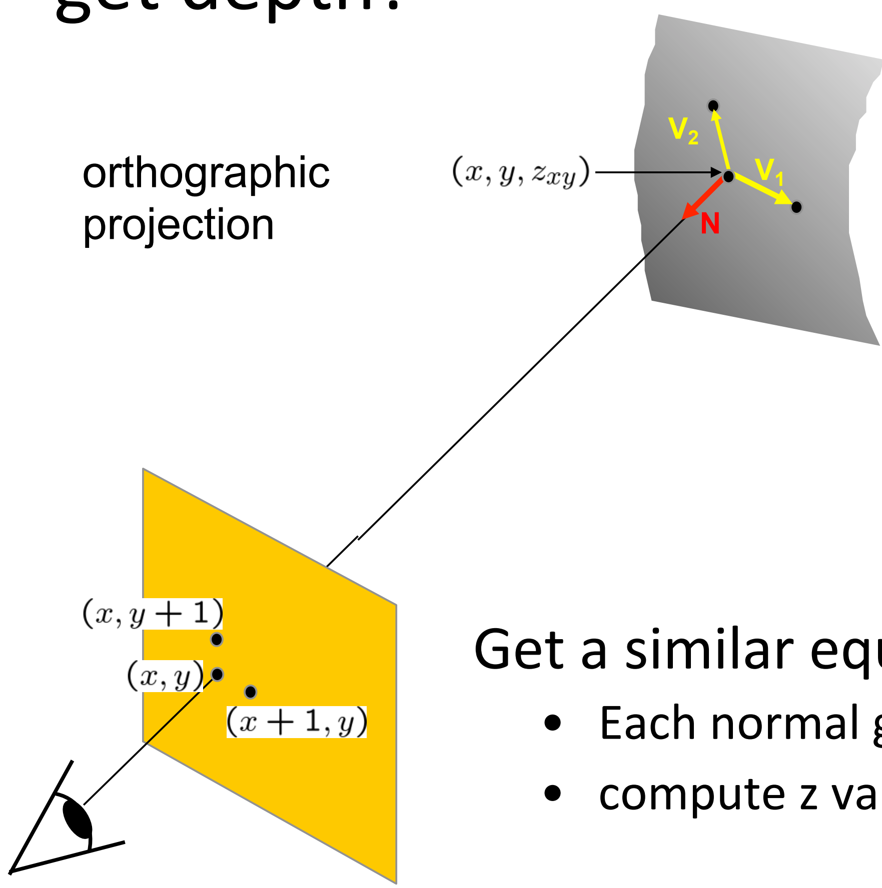

Assume a smooth surface

$$
\begin{aligned}
V_1 &= (x+1, y, z_{x+1,y}) - (x, y, z_{xy}) \\
&= (1, 0, z_{x+1,y} - z_{xy})
\end{aligned}
$$

$$
\begin{aligned}
0 &= N \cdot V_1 \\
&= (n_x, n_y, n_z) \cdot (1, 0, z_{x+1,y} - z_{xy}) \\
&= n_x + n_z(z_{x+1,y} - z_{xy})
\end{aligned}
$$

Get a similar equation for $\mathbf{V_2}$

- Each normal gives us two linear constraints on z
- compute z values by solving a matrix equation

Determining Light Directions

- Trick: Place a mirror ball in the scene.

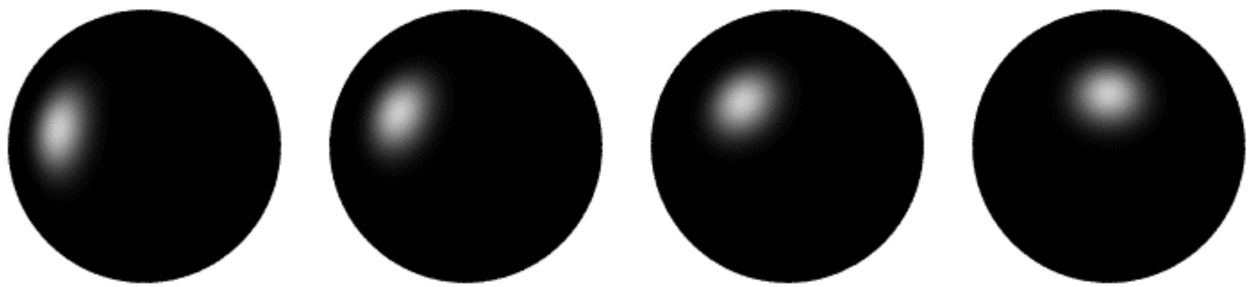

- The location of the highlight is determined by the light source direction.

Real-World HDR Lighting Environments

Funsto
Beach

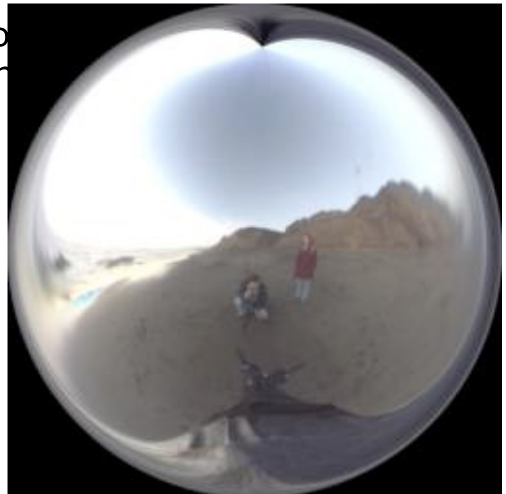

Eucalyptu
Grove

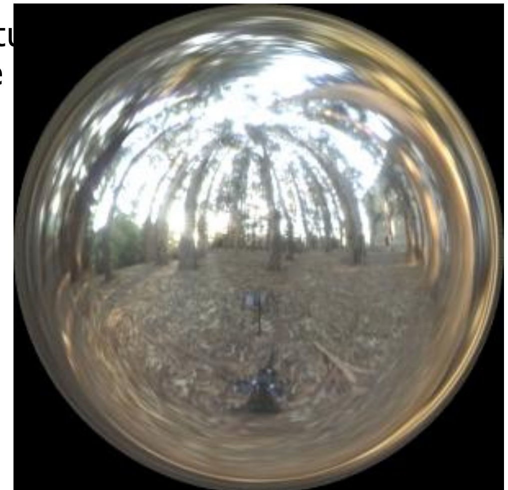

Uffizi
Gallery

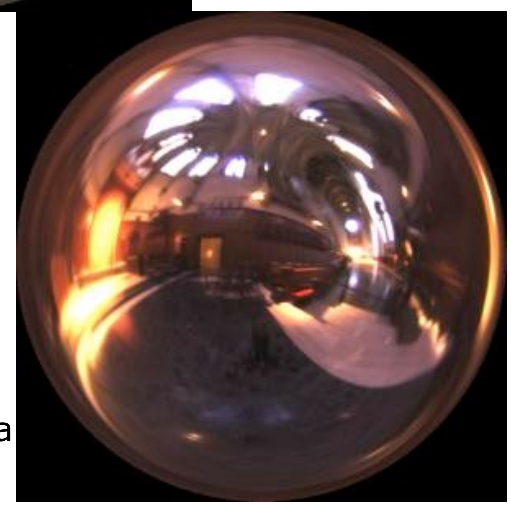

Grace
Cathedra

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


Mirrored Sphere

Canon
Canon
REMOTE SWITCH

Acquiring the Light Probe

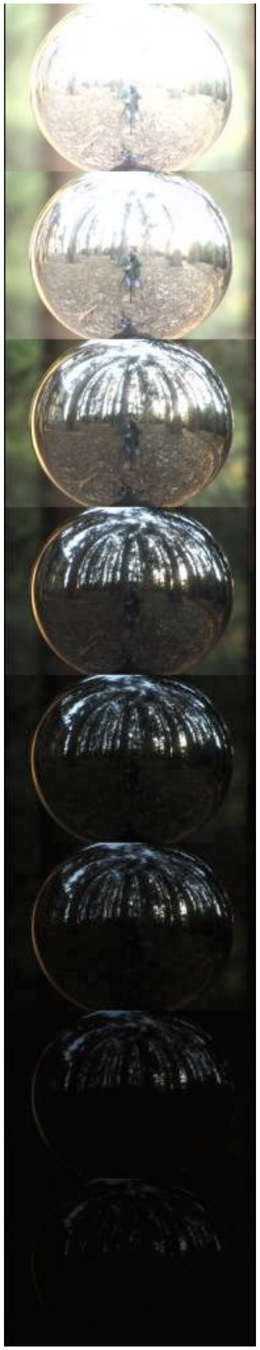

Assembling the Light Probe

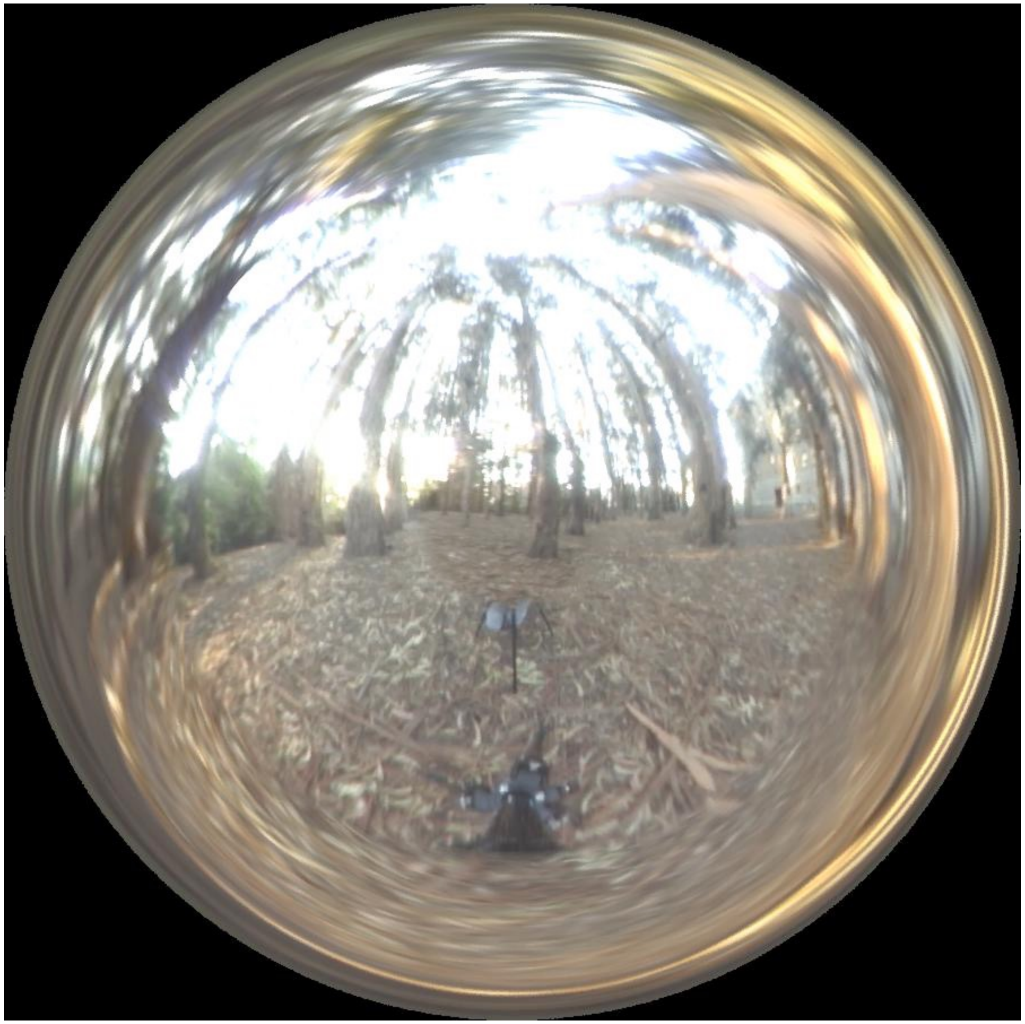

Extreme HDR Image Series

1 sec
f/4

1/4 sec
f/4

1/30 sec
f/4

1/30 sec
f/16

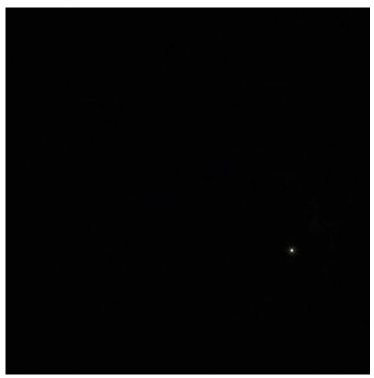

1/250 sec
f/16

1/1000 sec
f/16

1/8000 sec f/16

Extreme HDR Image Series

sun closeup

1 sec
f/4

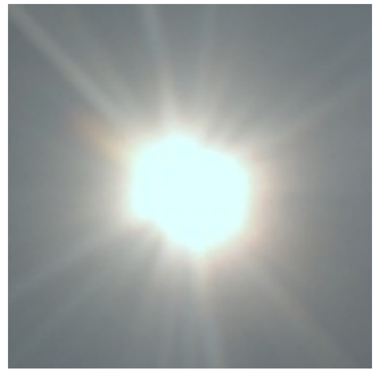

1/4 sec
f/4

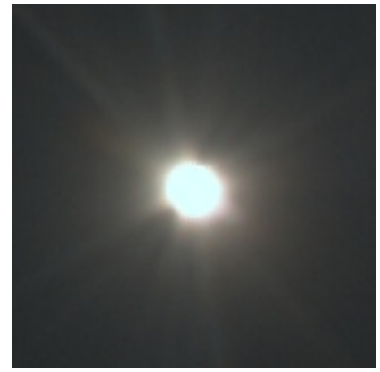

1/30 sec
f/4

1/30 sec
f/16

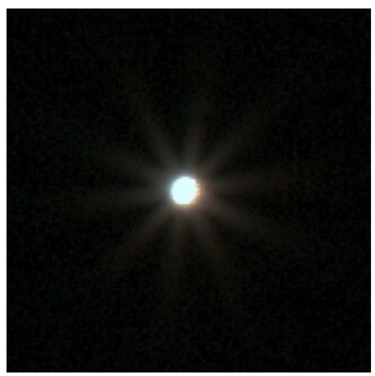

1/250 sec
f/16

1/1000 sec
f/16

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

HDRI Sky Probe

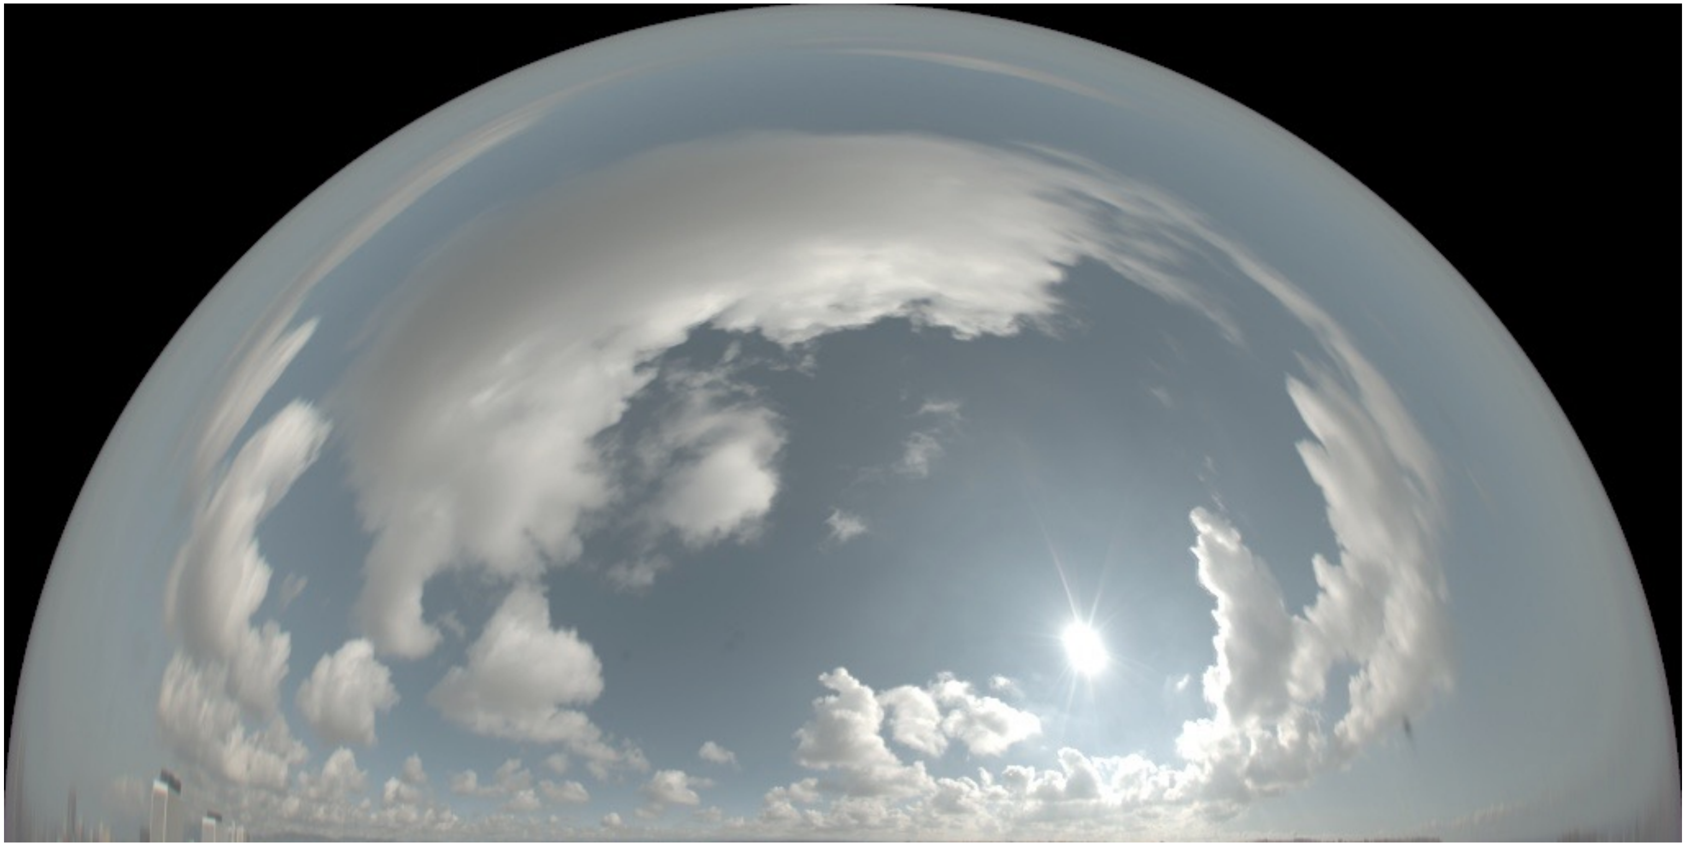

Determining Light Directions

- Trick: Place a mirror ball in the scene.

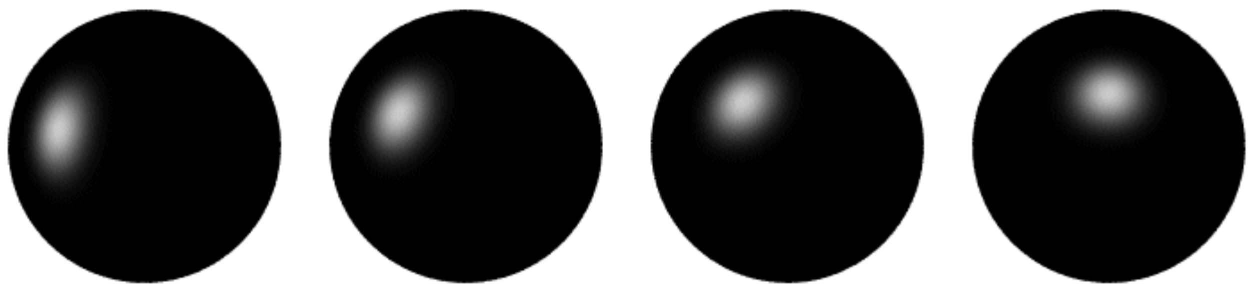

- The location of the highlight is determined by the light source direction.

Determining Light Directions

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

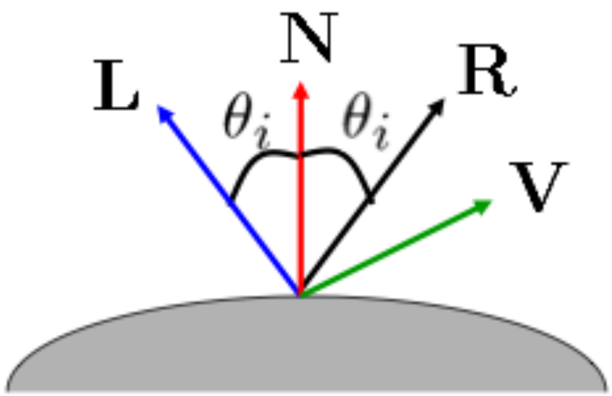

$$I_e = \begin{cases} I_i & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

- So the light source direction is given by:

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

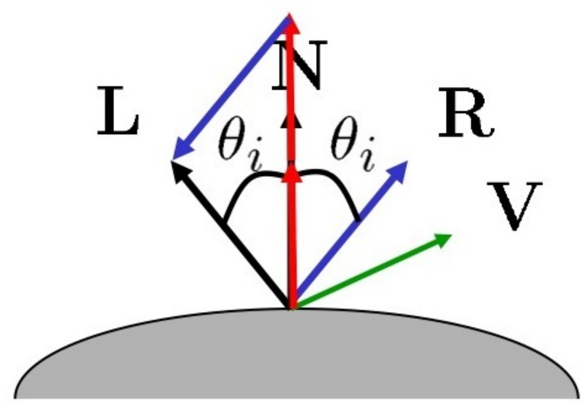

Determining Light Directions

- For a sphere with highlight at point H:

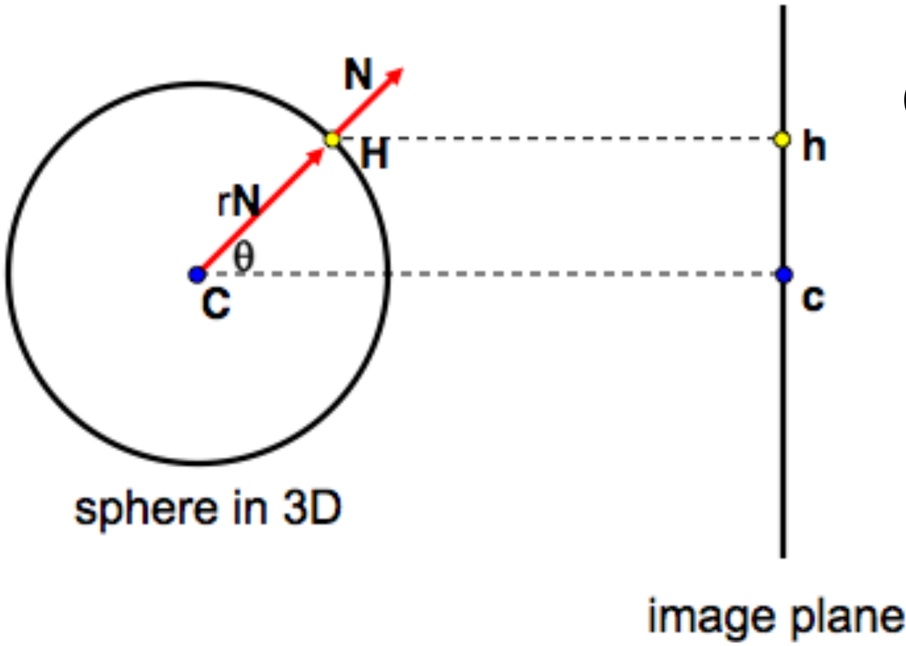

Compute N:

$$N_x = \frac{x_h - x_c}{r}$$

$$N_y = \frac{y_h - y_c}{r}$$

$$N_z = \sqrt{1 - x^2 - y^2}$$

- R = direction of the camera from C = $[0\ 0\ 1]^T$

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

Results

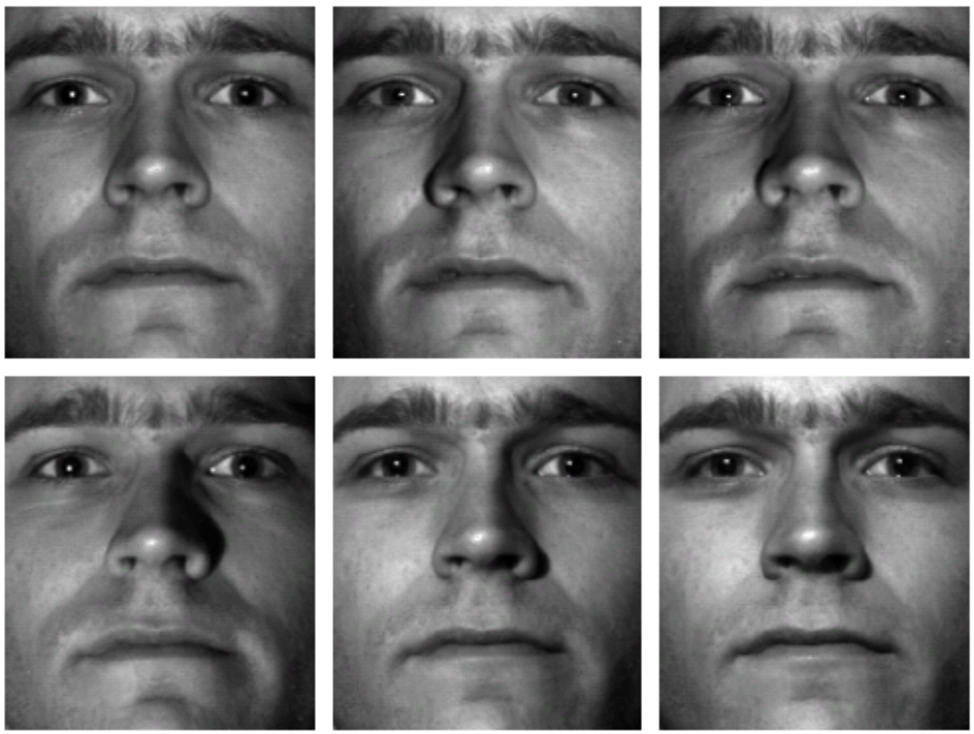

from Athos Georghiades

Results

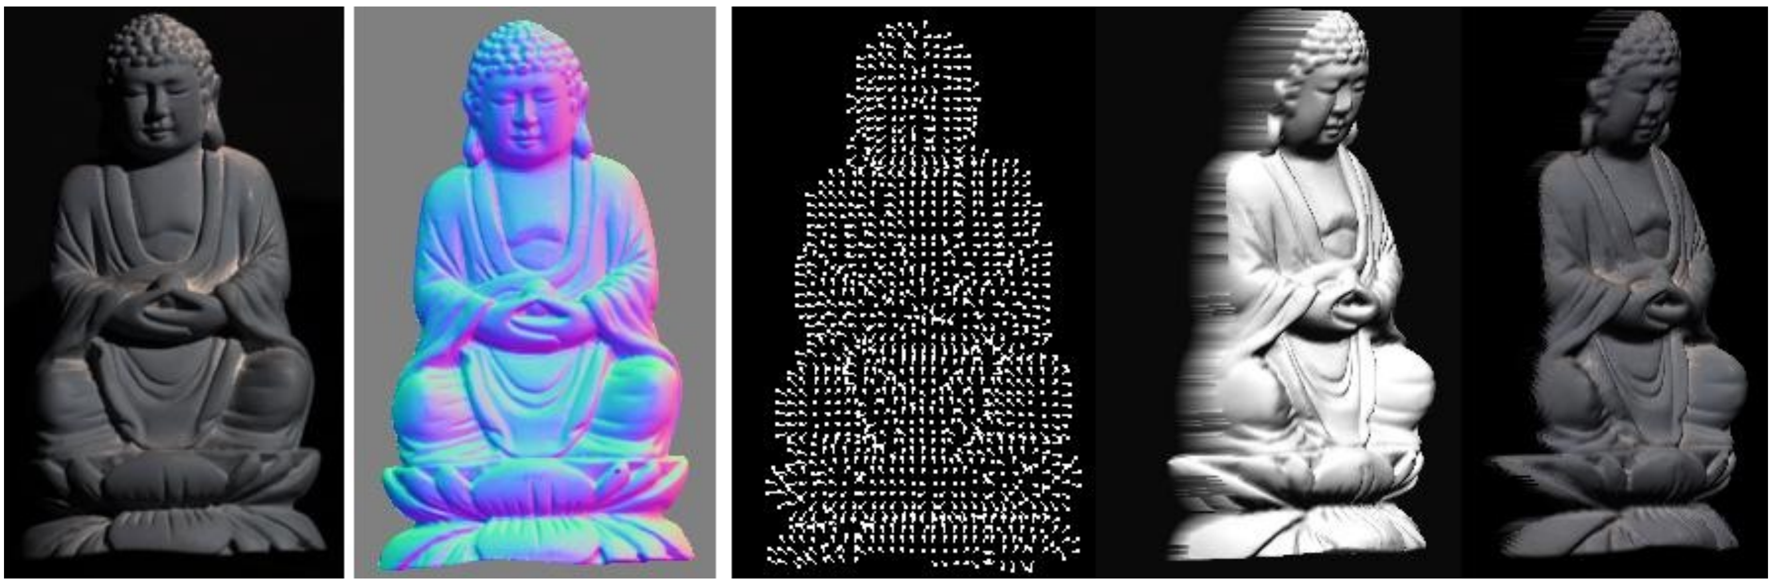

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Color Images

- Now we have 3 equations for a pixel:

$$\begin{aligned} \mathbf{I}_R &= k_{dR}\,\mathbf{LN} \\ \mathbf{I}_G &= k_{dG}\,\mathbf{LN} \\ \mathbf{I}_B &= k_{dB}\,\mathbf{LN} \end{aligned}$$

- Simple approach: solve for N using grayscale or a single channel
- Then fix N and solve for each channel's : k_d

$$k_d = \frac{\sum_i I_i L_i N^T}{\sum_i (L_i N^T)^2}$$

Color Images

- Then fix N and solve for each channel's : k_d

$$Q = \sum_i (I_i - k_d L_i N^T)^2$$

$$\frac{\delta Q}{\delta k_d} = \sum_i -2(I_i - k_d L_i N^T) L_i N^T = 0$$

$$k_d = \frac{\sum_i I_i L_i N^T}{\sum_i (L_i N^T)^2}$$

For (unfair) Comparison

- Multi-view stereo results on a similar object
- 47+ hrs compute time

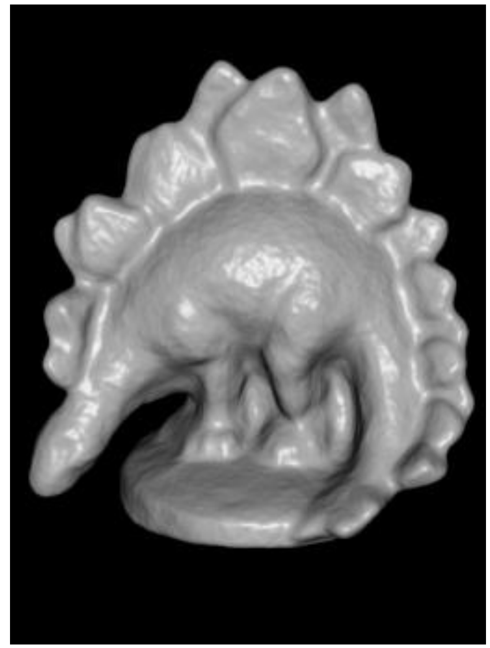
State-of-the-art MVS result


Ground truth

Taking Stock: Assumptions

Lighting	Materials	Geometry	Camera
directional	diffuse	convex / no shadows	linear
known direction	no inter-reflections		orthographic
> 2 nonplanar directions	no subsurface scattering		

Questions?

Unknown Lighting

- What we' ve seen so far: [Woodham 1980]
- Next up: Unknown light directions [Hayakawa 1994]

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

Same as before, just transposed:

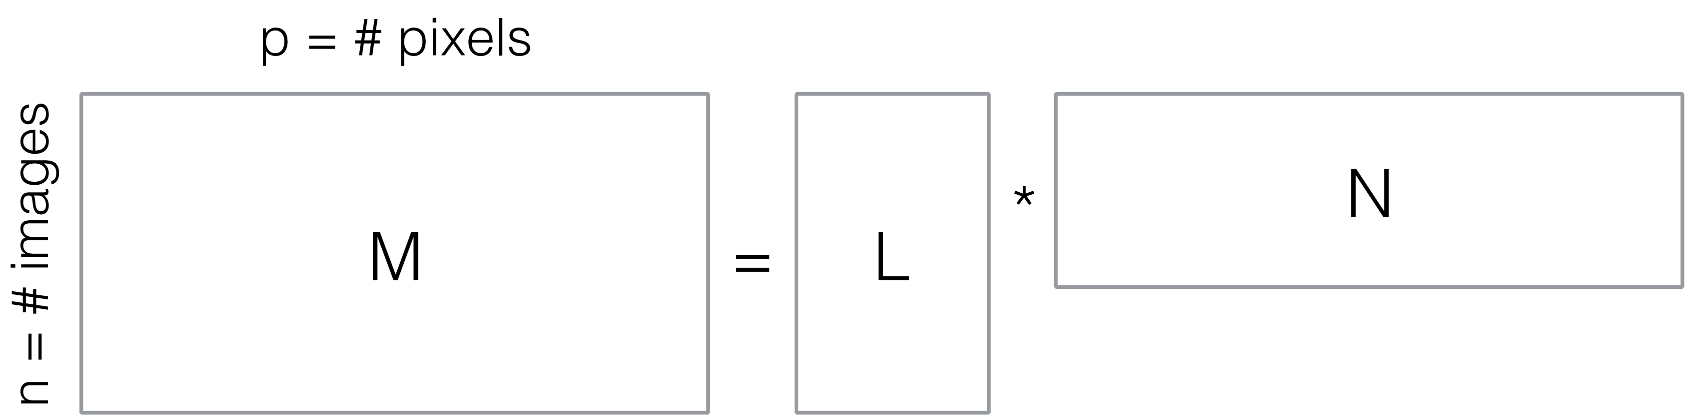

Unknown Lighting

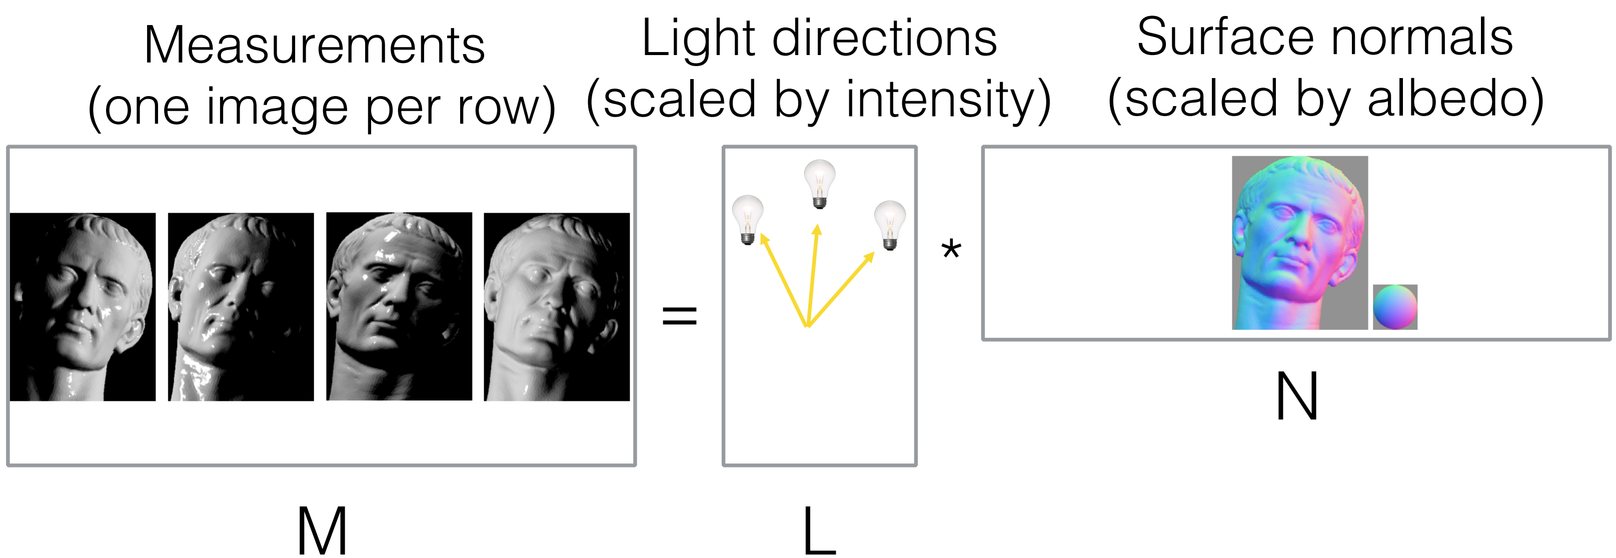

Both L and N are now unknown!
This is a matrix factorization problem.

Unknown Lighting

$$M_{ij} = L_i \cdot N_j$$

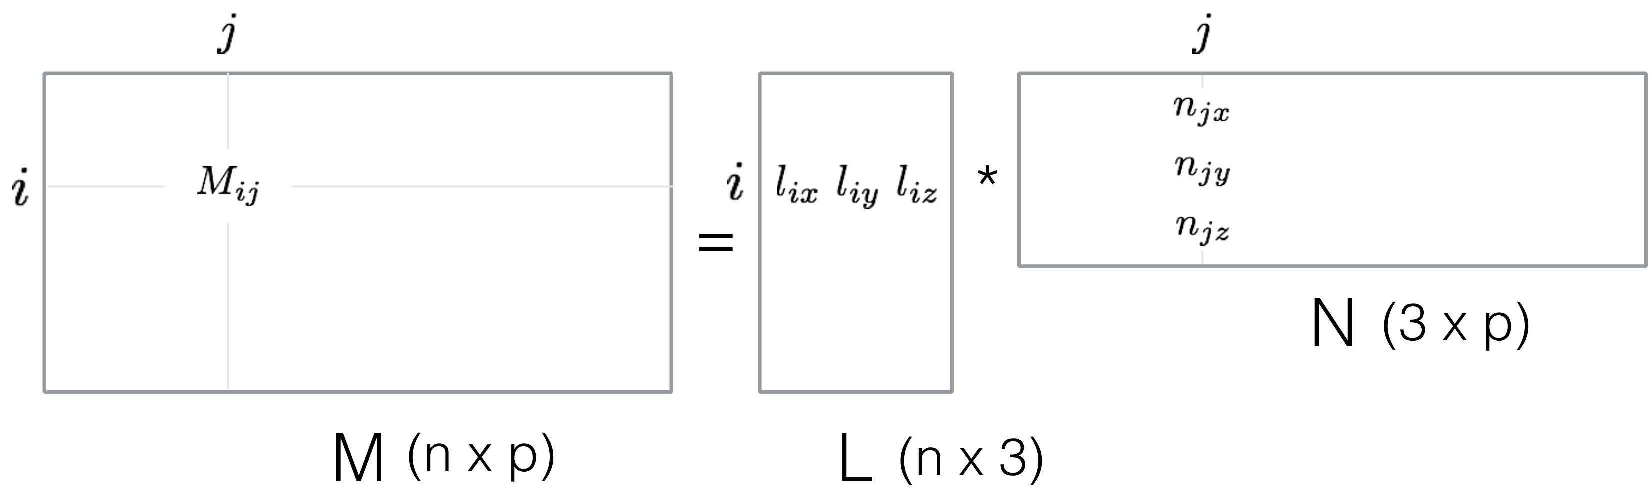

There's hope: We know that M is rank 3

Unknown Lighting

Use the SVD to decompose M:

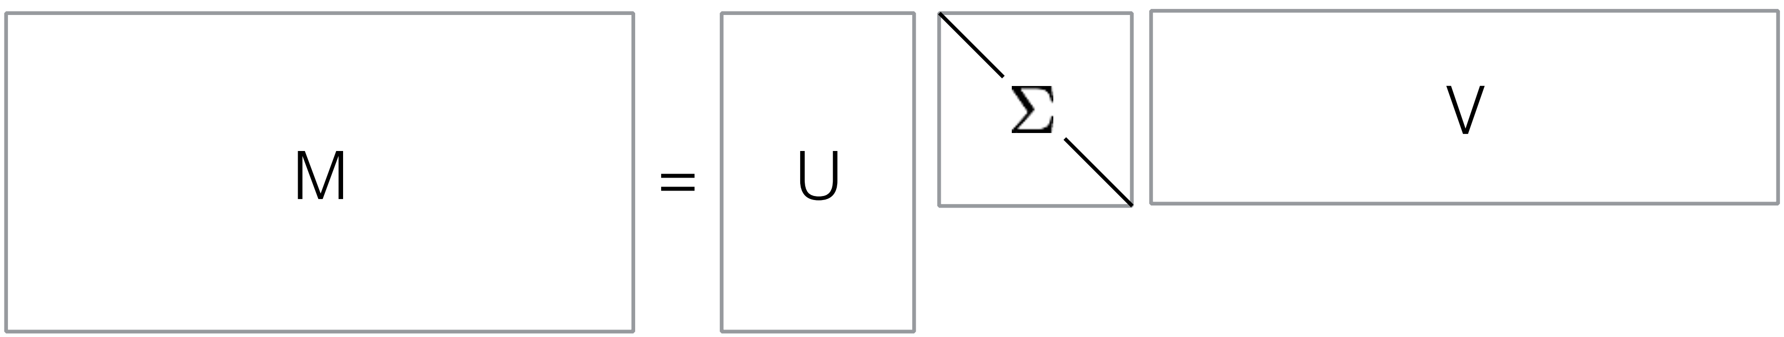

SVD gives the best rank-3 approximation of a matrix.

Unknown Lighting

Use the SVD to decompose M:

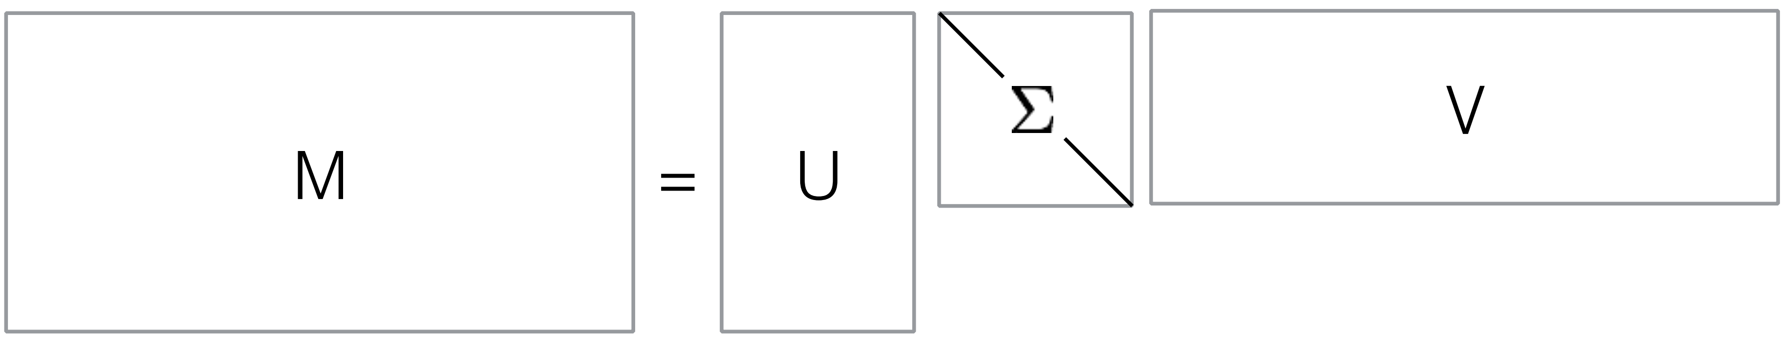

What do we do with Σ?

Unknown Lighting

Use the SVD to decompose M:

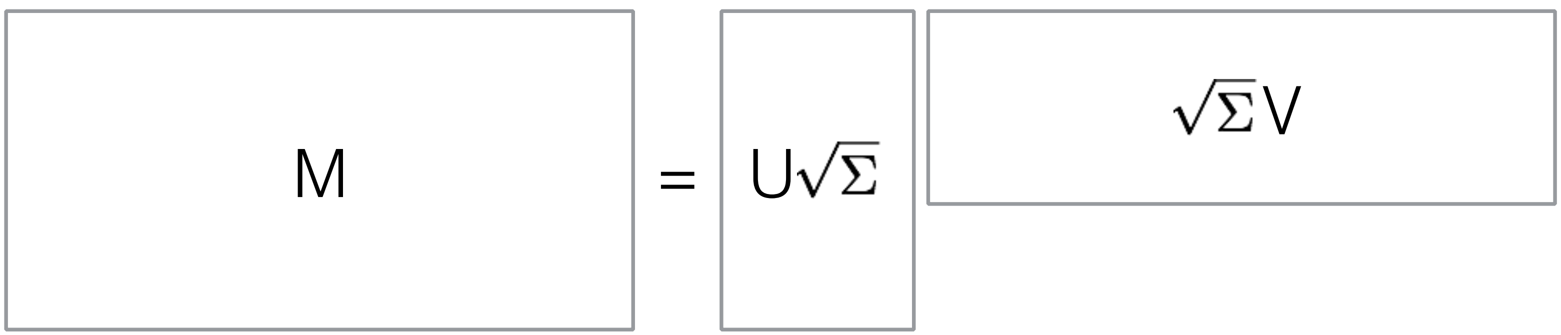

What do we do with Σ?

Unknown Lighting

Use the SVD to decompose M:

$$M = U\sqrt{\Sigma}\;\sqrt{\Sigma}V$$

Can we just do that?

Unknown Lighting

Use the SVD to decompose M:

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

$$\hat{L} = U\sqrt{\Sigma}, \hat{S} = \sqrt{\Sigma}V$$

Can we just do that? …almost.

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

Unknown Lighting

Use the SVD to decompose M:

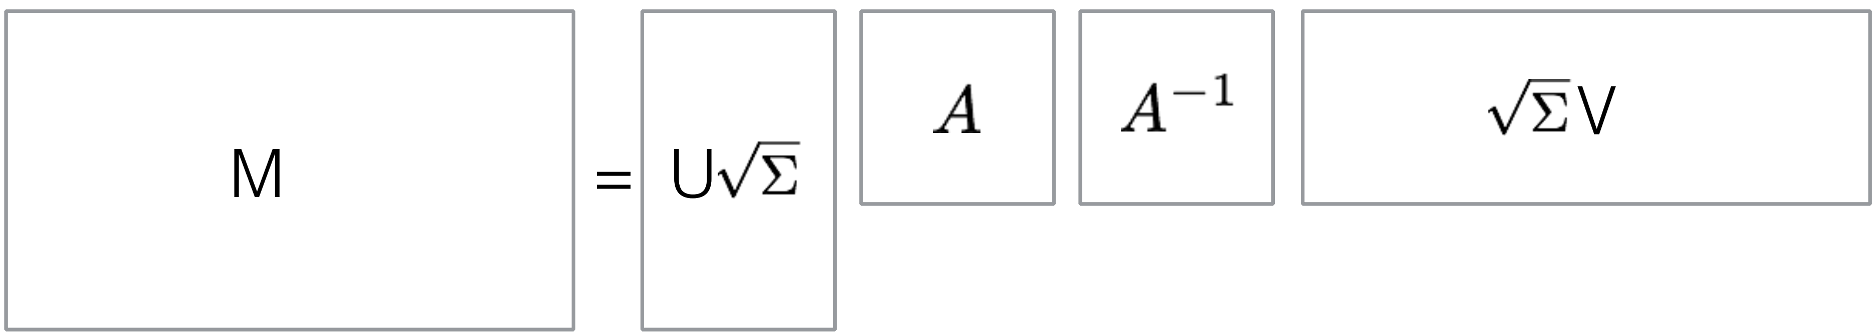

$$L = U\sqrt{\Sigma}A \qquad S = A^{-1}\sqrt{\Sigma}V$$

Unknown Lighting

Use the SVD to decompose M:

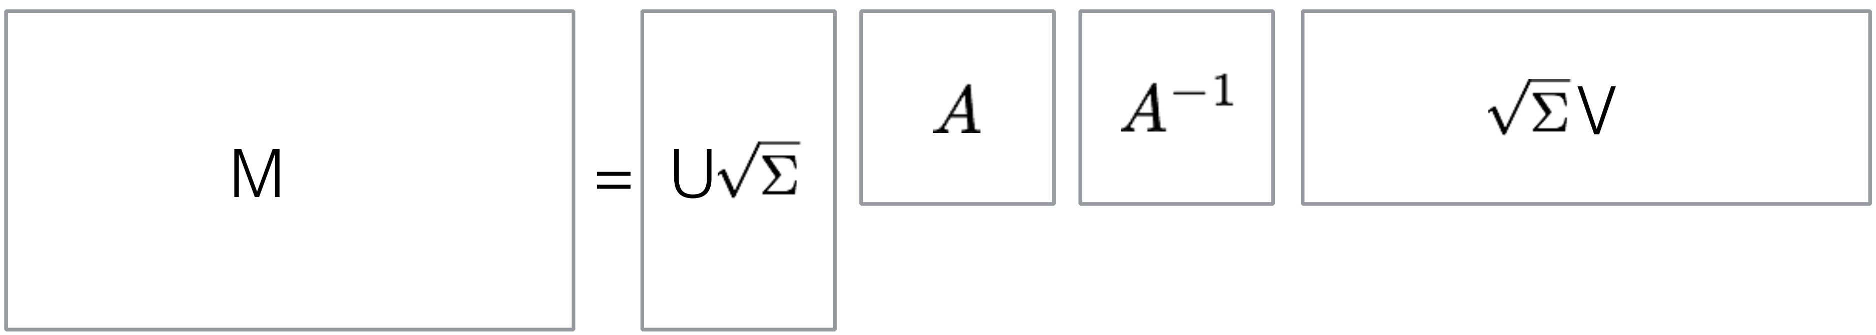

$$L = U\sqrt{\Sigma}A \qquad S = A^{-1}\sqrt{\Sigma}V$$

You can find A if you know

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

Why 6 points?

- Let C = A^{-1}, to match the notation of paper

$$\hat{S} = \sqrt{\Sigma} V$$

$$\hat{s_k}^T C C^T \hat{s_k} = constant$$

$$\hat{s_k}^T B \hat{s_k} = constant$$

- B is symmetric, hence 6

$$B = W \Pi W^T, C = W \sqrt{\Pi}$$

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.

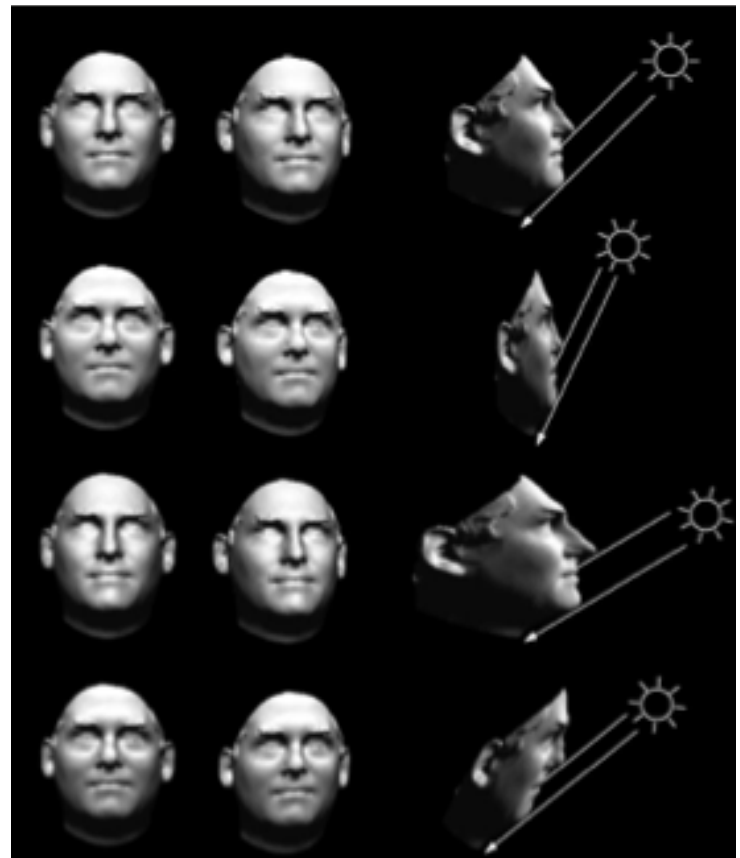

[Belhumeur et al.'97]

Since 1994…

- Workarounds for many of the restrictive assumptions.
- Webcam photometric stereo:

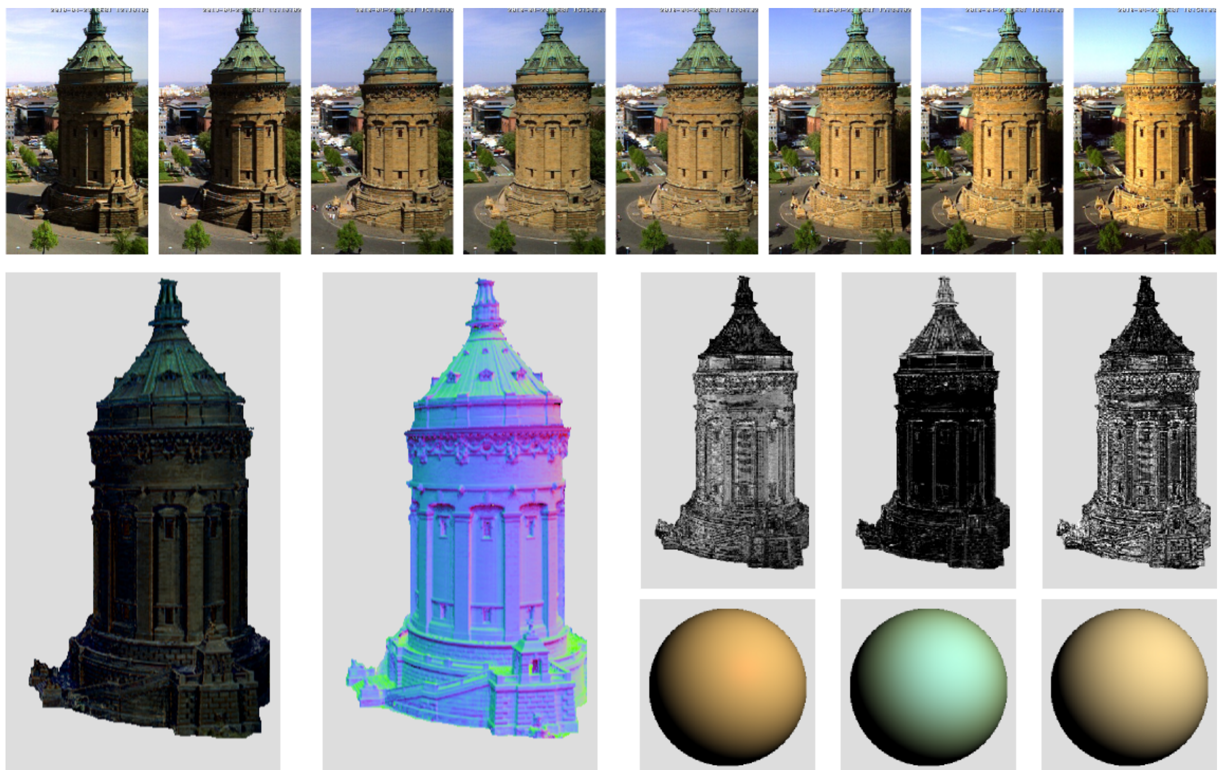

Ackermann et al. 2012

Since 1994…

- Photometric stereo from unstructured photo collections (different cameras and viewpoints):

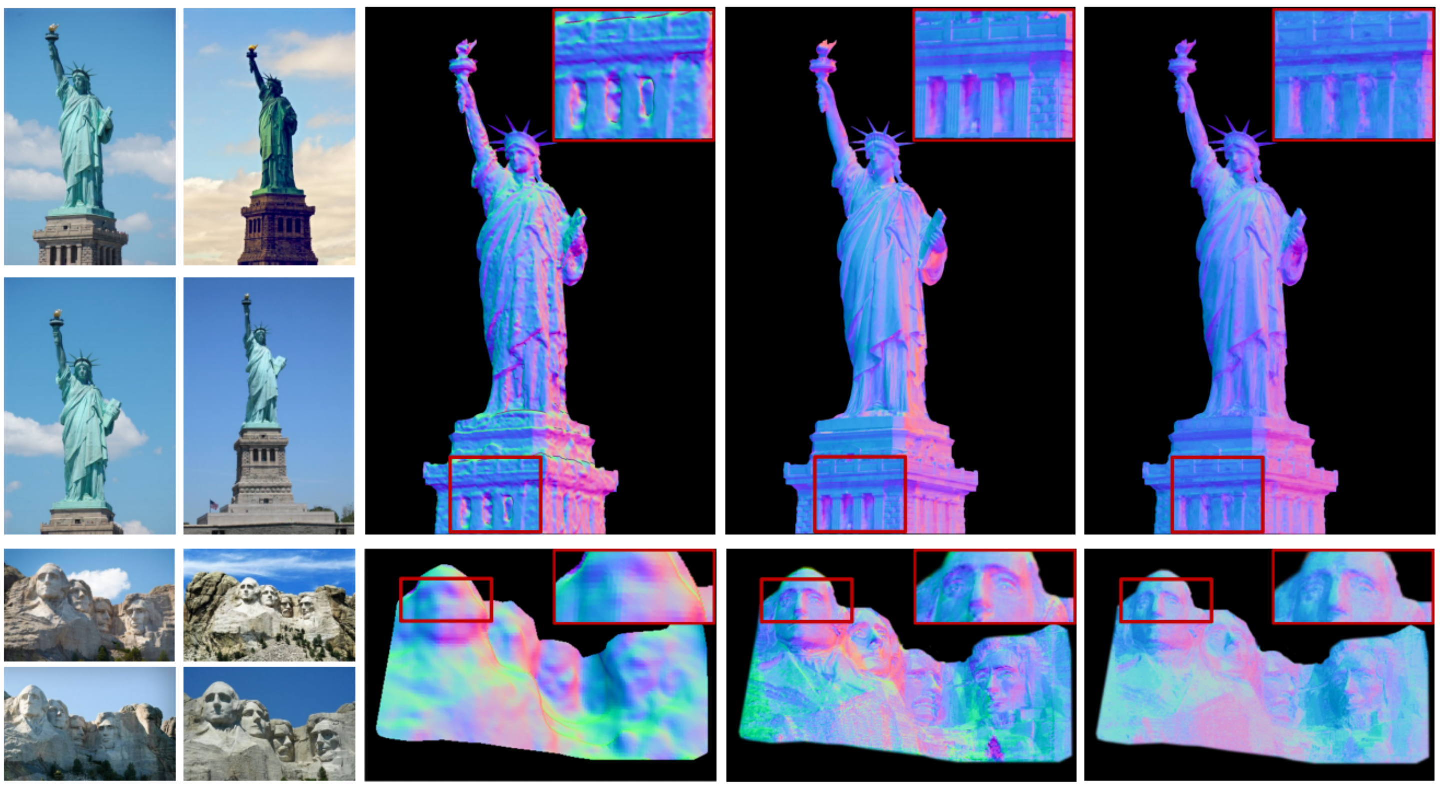

Shi et al, 2014

Since 1994…

- Non-Lambertian (shiny) materials:

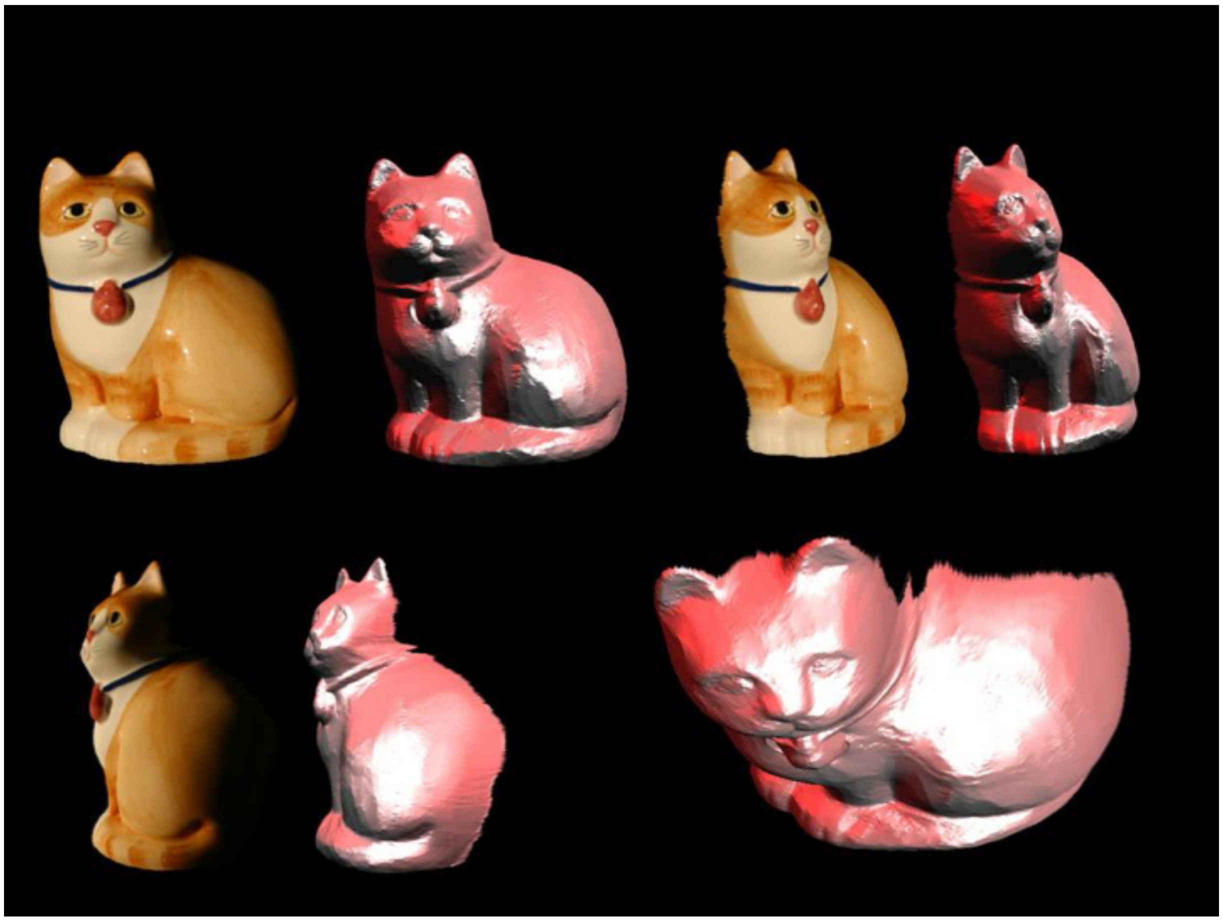

Hertzmann and Seitz, 2005

Since 1994...

- Local visibility

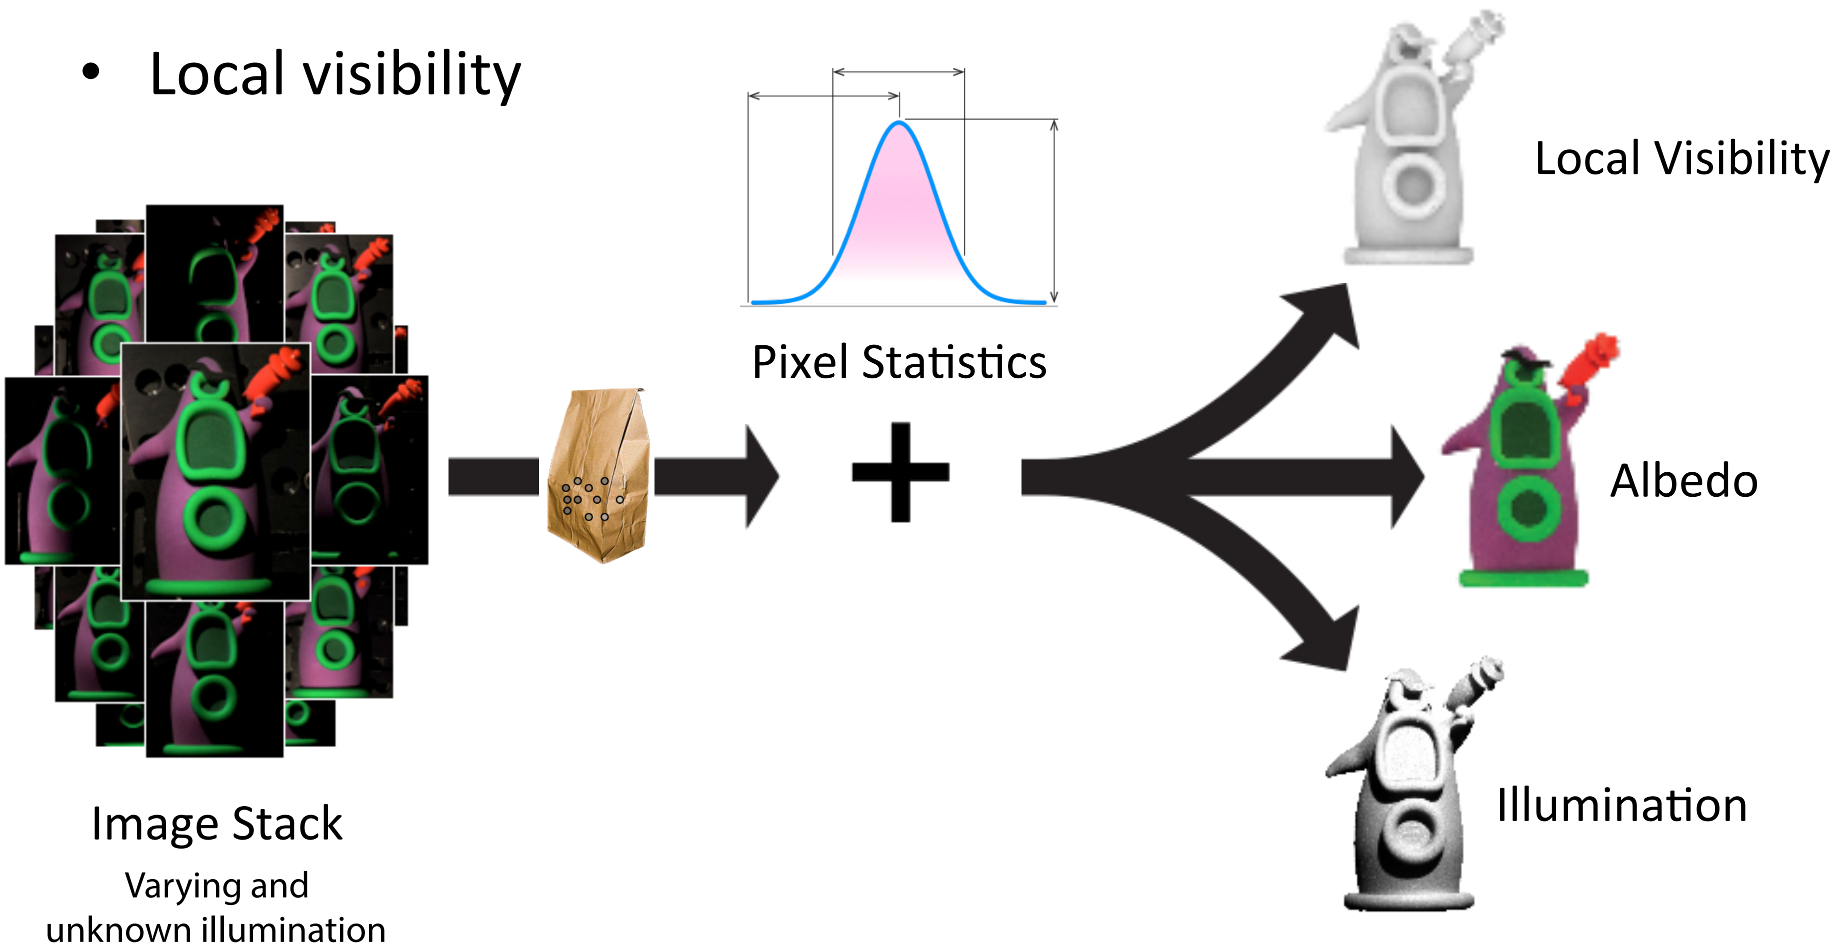

Hauagge, Wehrwein, Bala, Snavely, 2013

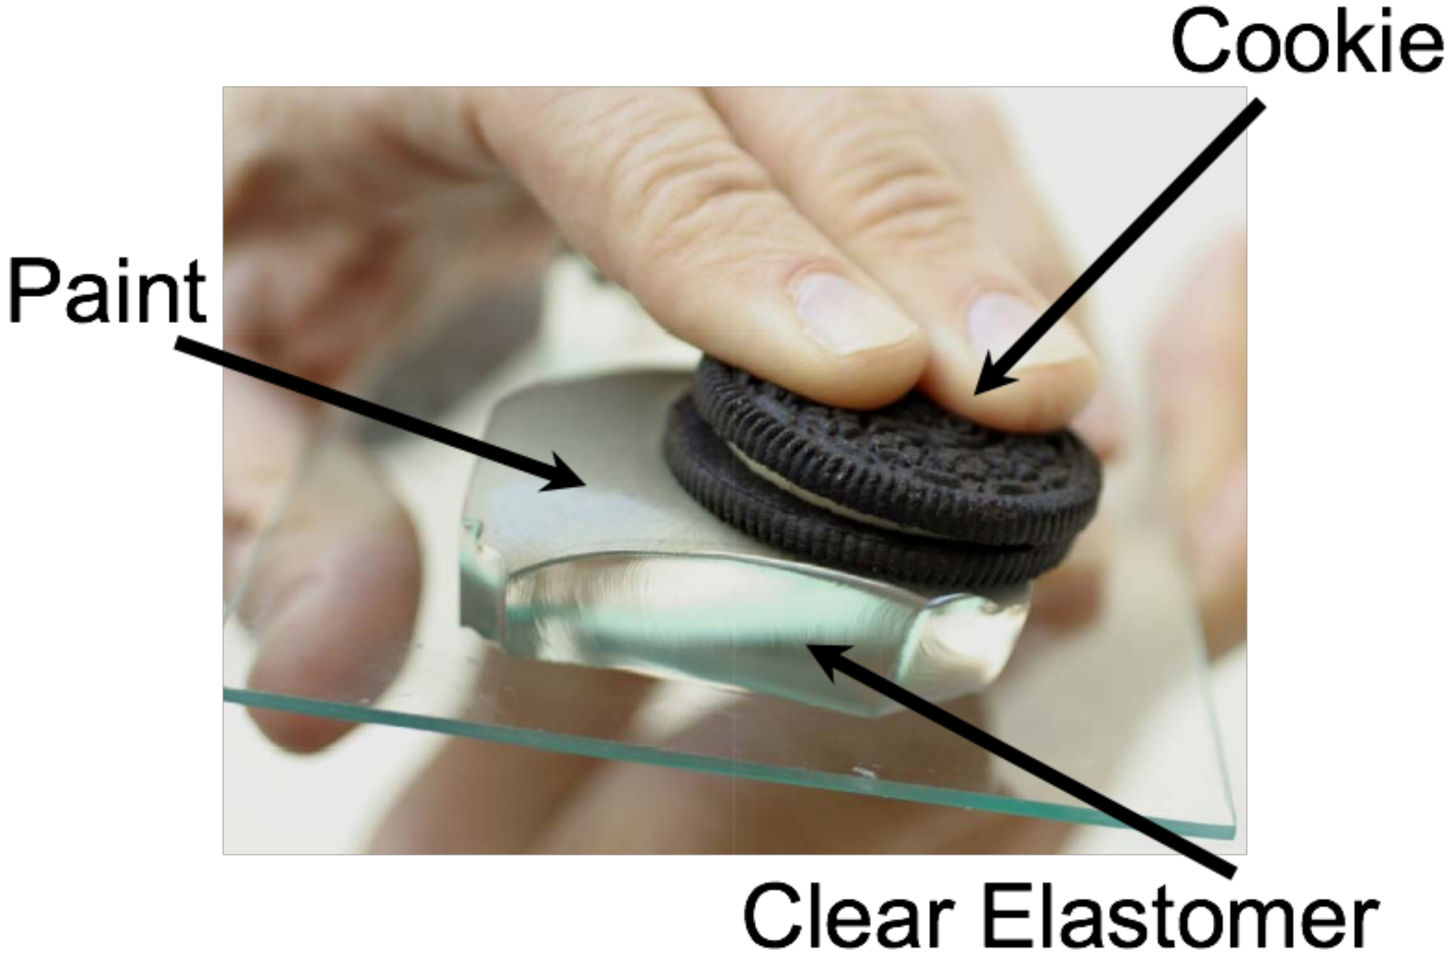

Johnson and Adelson, 2009


OREO


OREO

Lights, camera, action

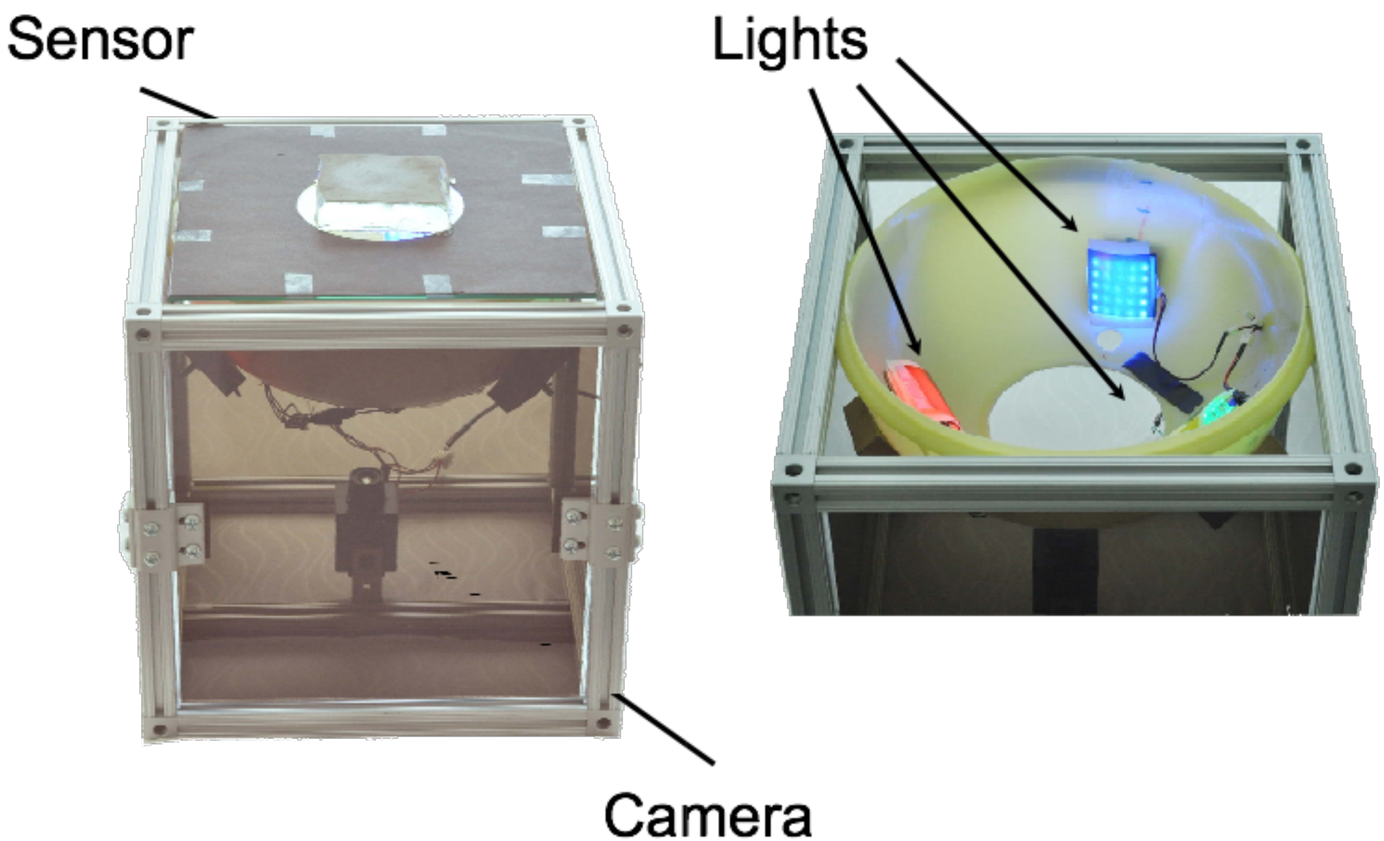

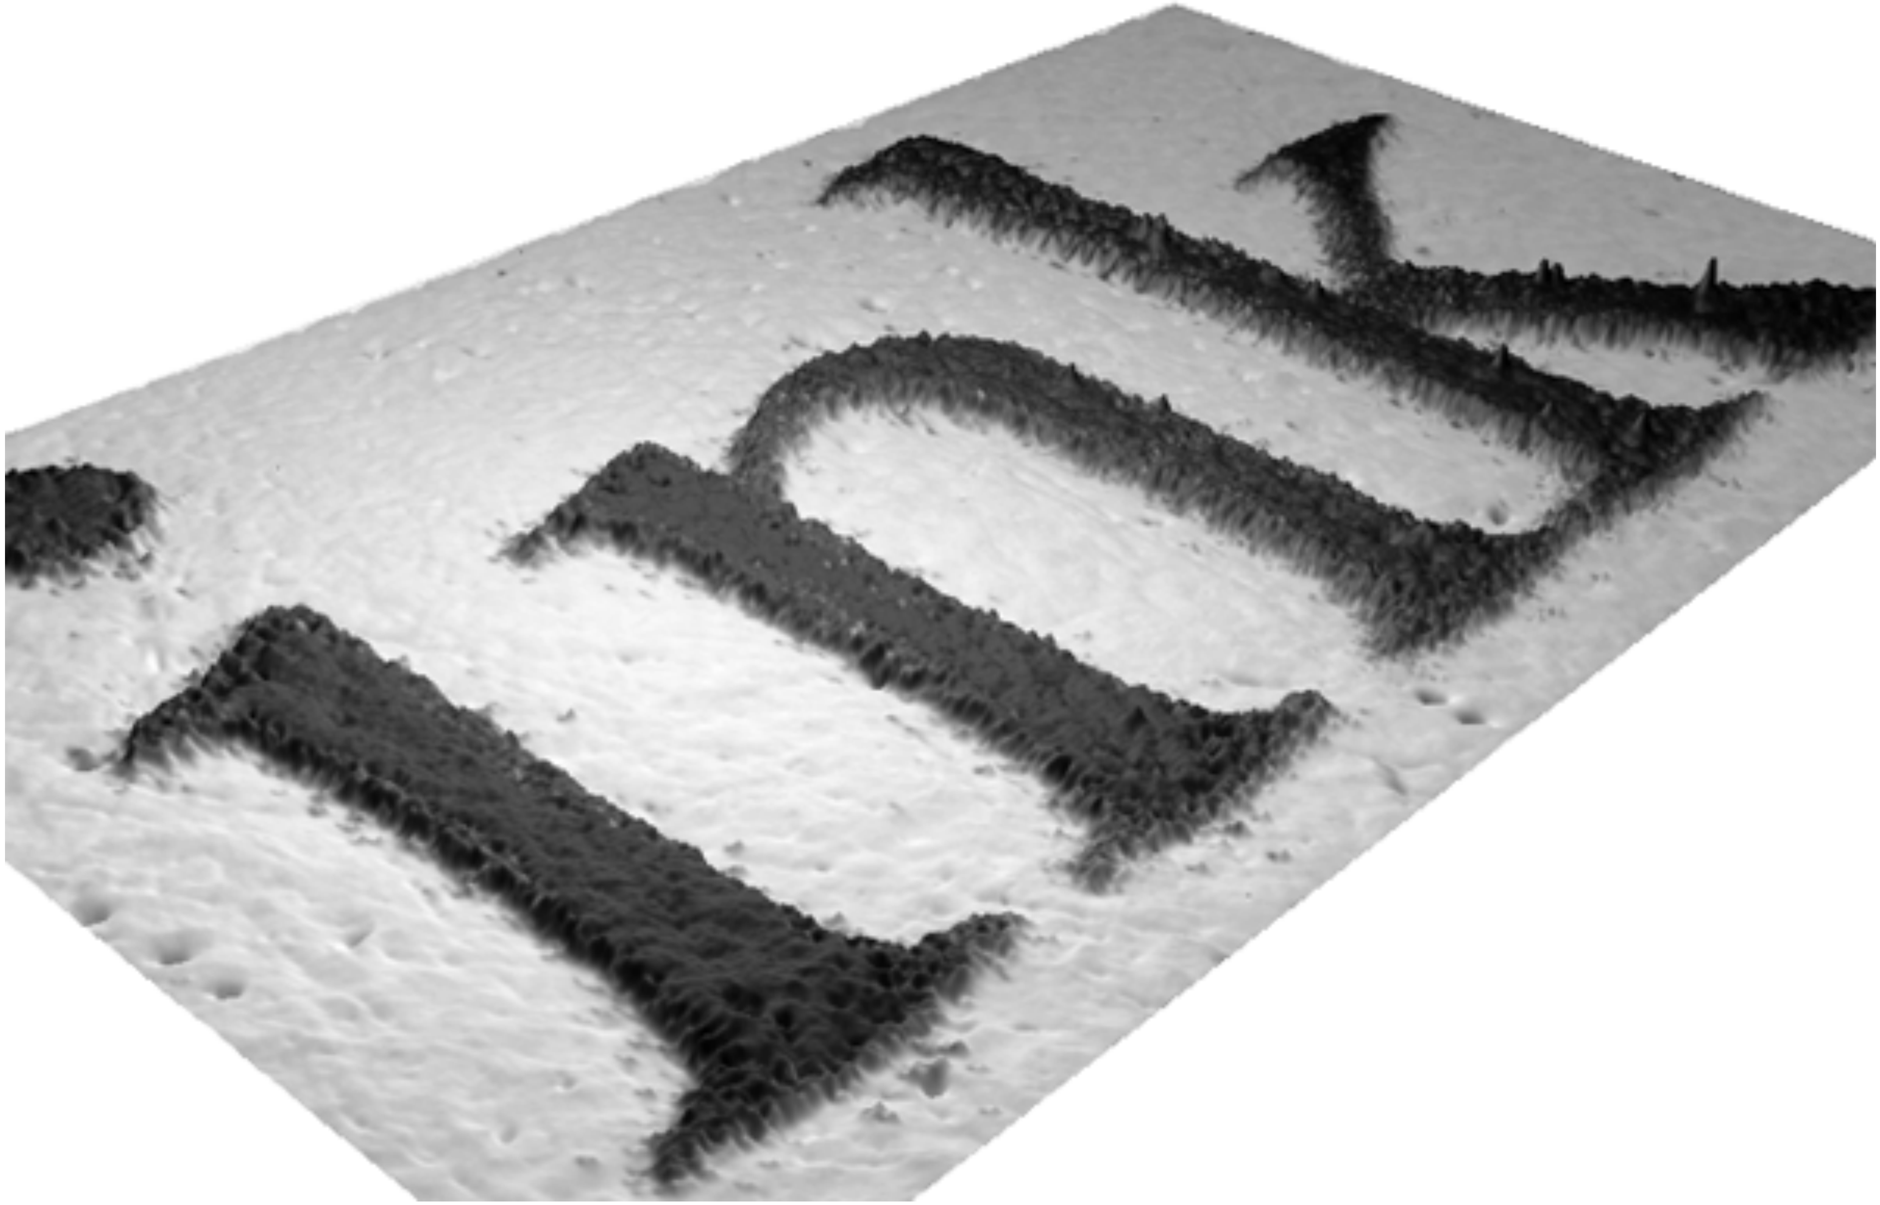

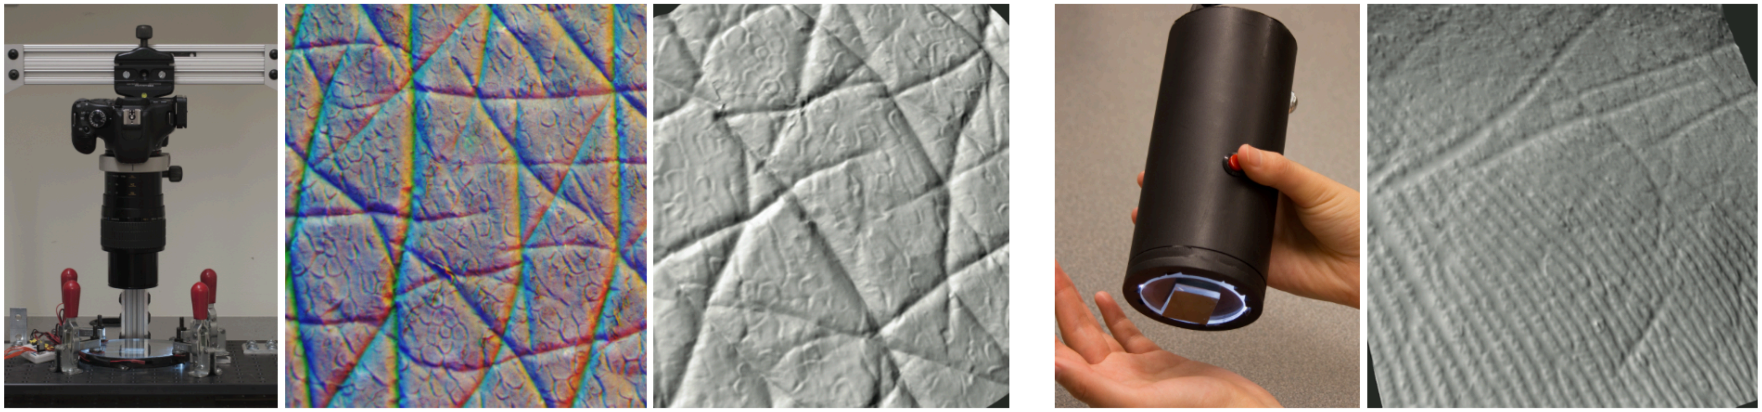

(a) bench configuration (b) captured (c) reconstruction (d) portable configuration (e) reconstruction

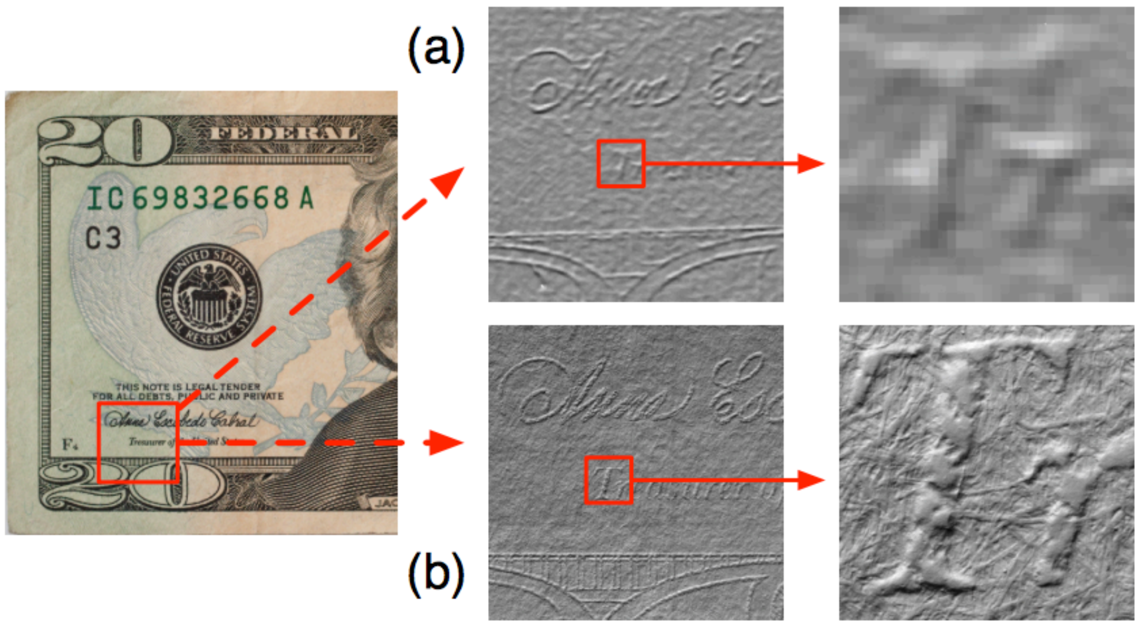

Figure 7: *Comparison with the high-resolution result from the original retrographic sensor. (a) Rendering of the high-resolution \$20 bill example from the original retrographic sensor with a close-up view. (b) Rendering of the captured geometry using our method.*

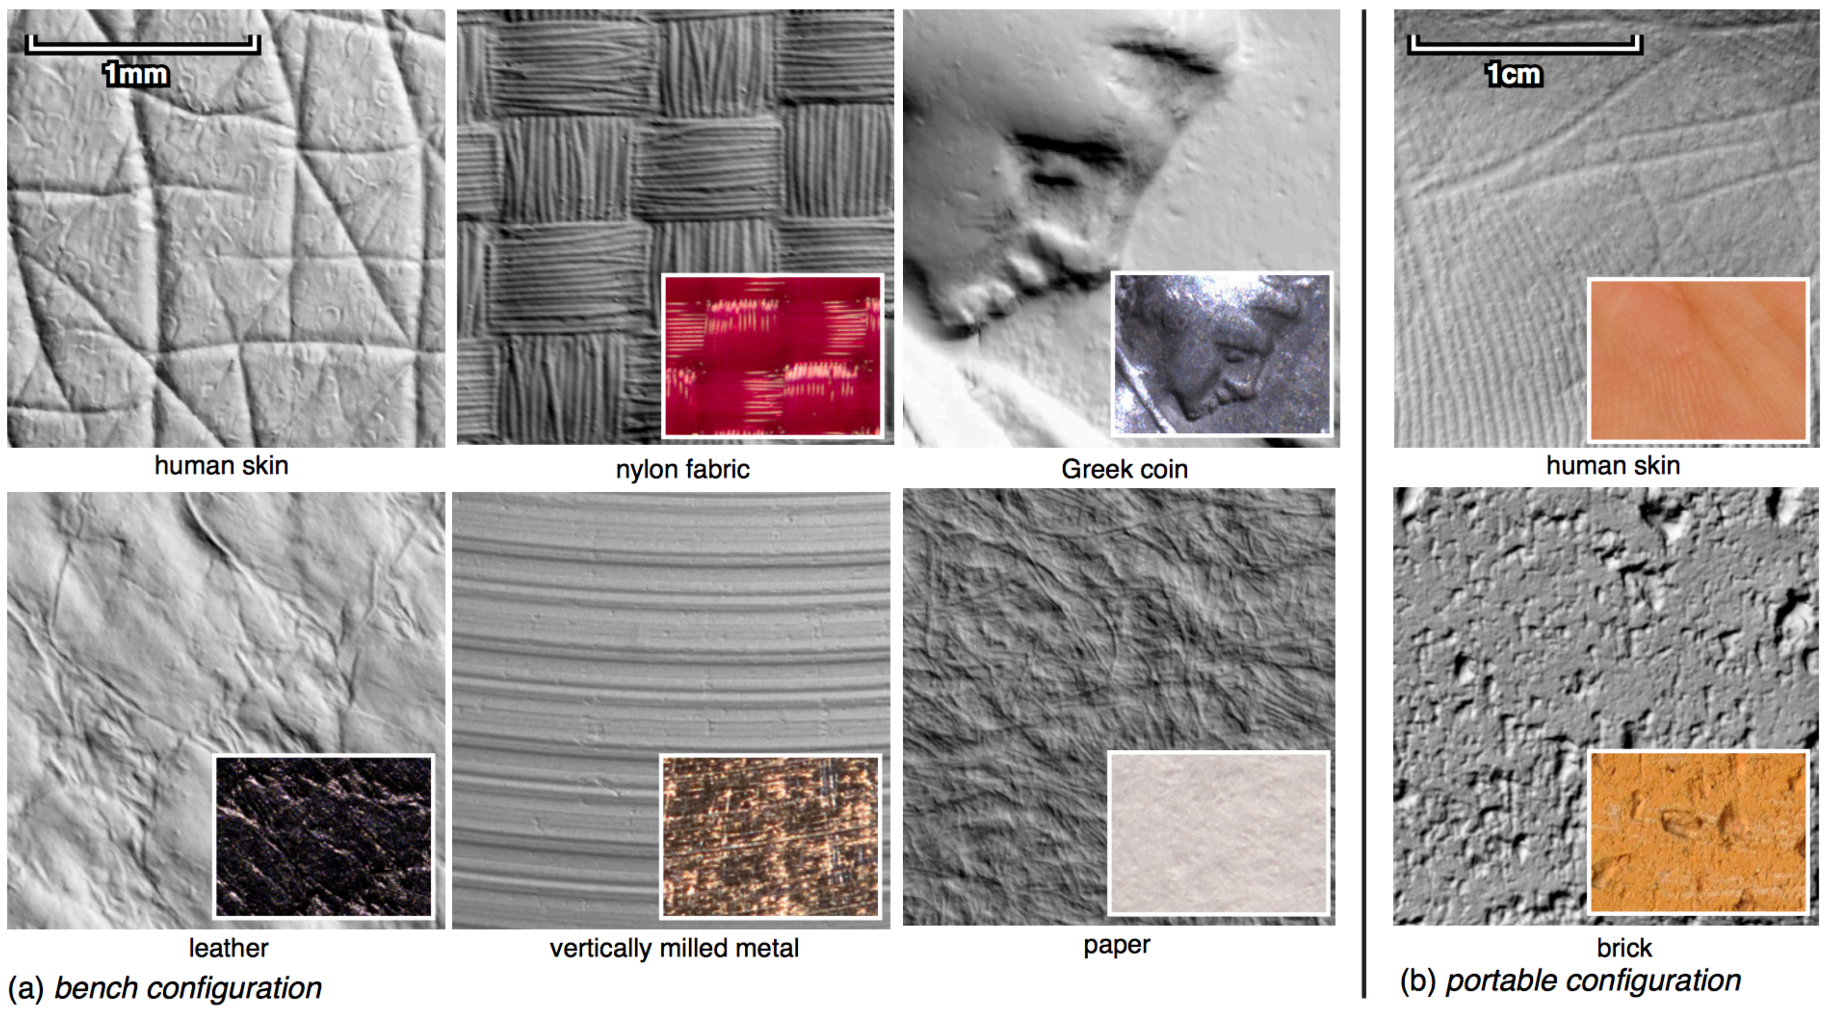

Figure 9: *Example geometry measured with the bench and portable configurations. Outer image: rendering under direct lighting. Inset: macro photograph of original sample. Scale shown in upper left. Color images are shown for context and are to similar, but not exact scale.*