Photometric stereo

Multiple Images: Photometric Stereo

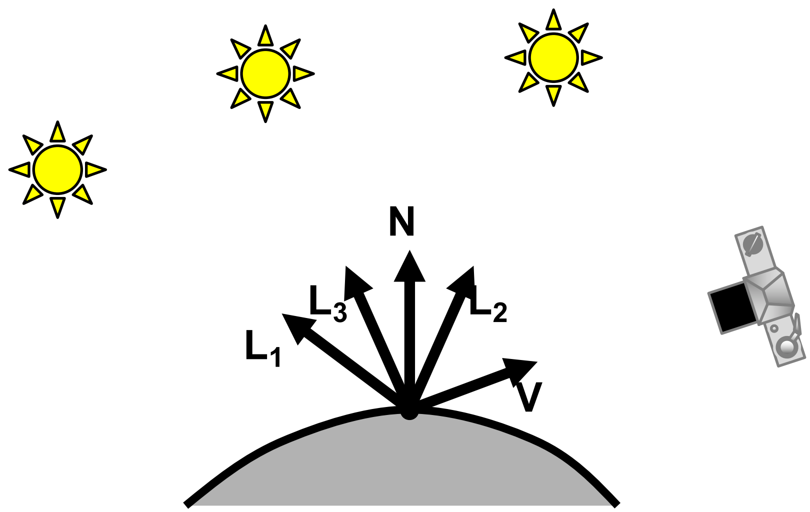

Photometric stereo

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

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

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

Multiple pixels: matrix form

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

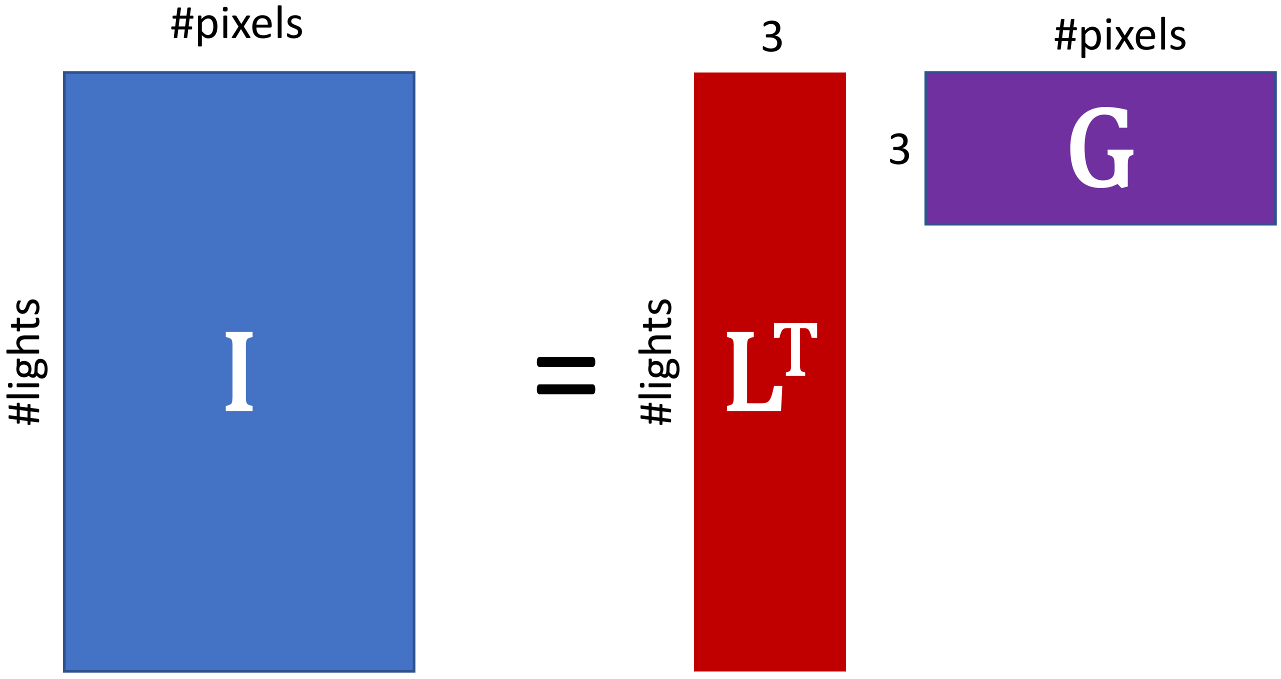

Normal equations

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

- 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 depth from normals

- So we got surface normals, can we get depth?
- Yes, given *boundary conditions*
- Normals provide information about the derivative

Brief detour: Orthographic projection

- Perspective projection
  - $x = \frac{X}{Z}, y = \frac{Y}{Z}$
- If all points have similar depth
  - $Z \approx Z_0$
  - $x \approx \frac{X}{Z_0}, y \approx \frac{Y}{Z_0}$
  - $x \approx cX, y \approx cY$
- A scaled version of orthographic projection
  - $x = X, y = Y$

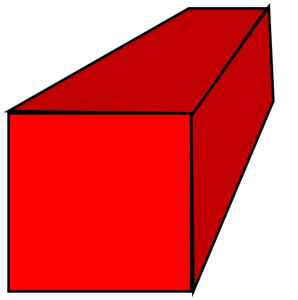

Perspective

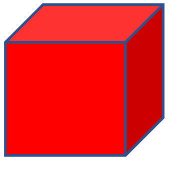

Scaled orthographic

Depth Map from Normal Map

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

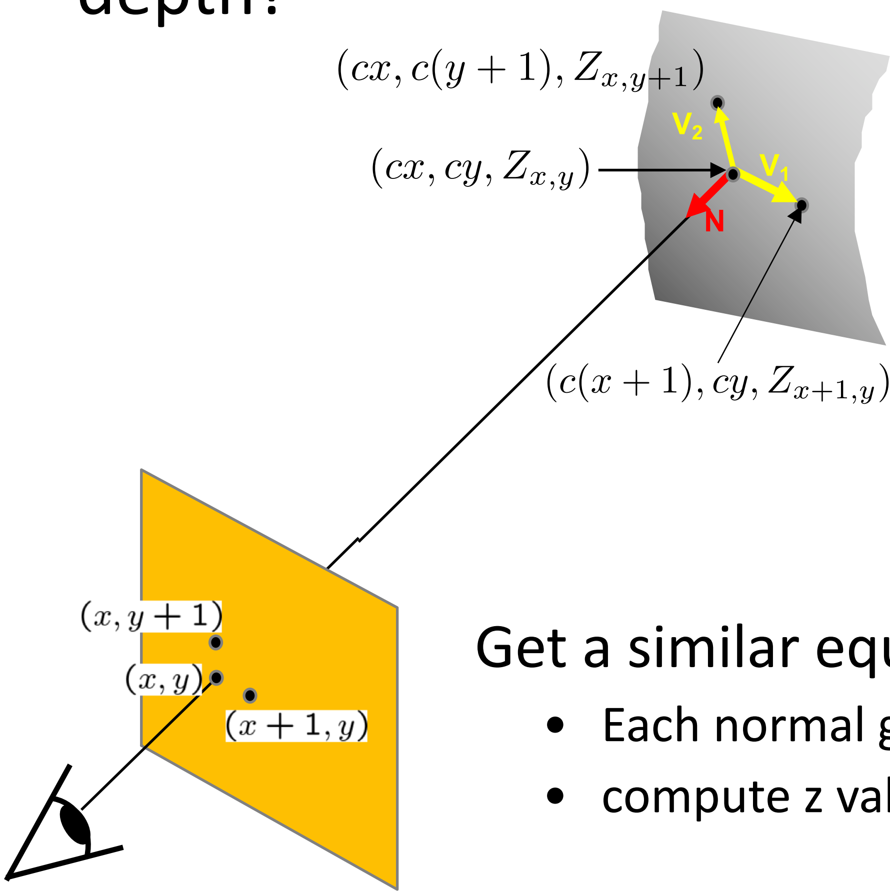

Assume a smooth surface

$$V_1 = (c(x+1), cy, Z_{x+1,y}) - (cx, cy, Z_{x,y})$$
$$= (c, 0, Z_{x+1,y} - Z_{x,y})$$

$$0 = N \cdot V_1$$
$$= (n_x, n_y, n_z) \cdot (c, 0, Z_{x+1,y} - Z_{x,y})$$
$$= cn_x + n_z(Z_{x+1,y} - Z_{x,y})$$

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.

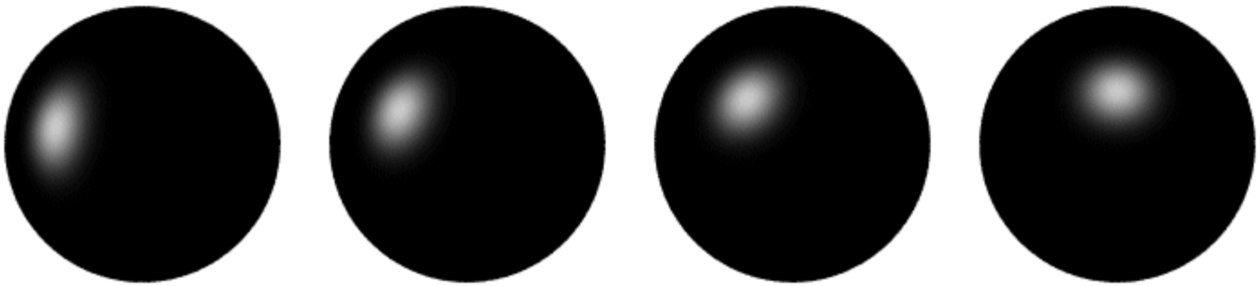

- The location of the highlight is determined by the light source direction.

Determining Light Directions

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

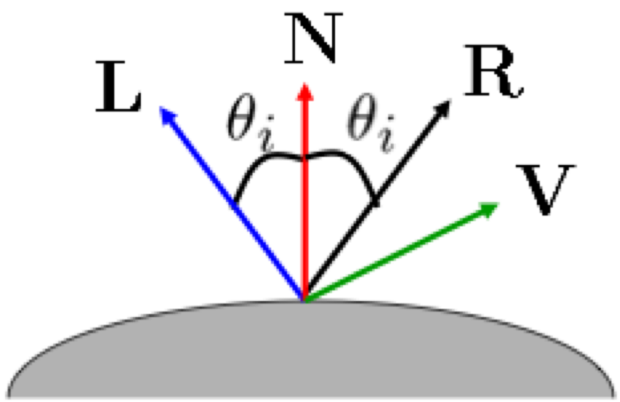

$$I_e = \begin{cases} I_i & \text{if } \mathbf{V} = \mathbf{R} \\ 0 & \text{otherwise} \end{cases}$$

Determining Light Directions

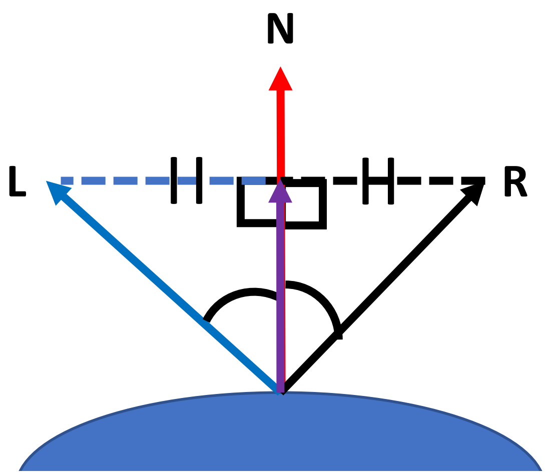

 $= (N \cdot R)N$

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

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

 =  - 2

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

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

So the light source direction is given by:

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

Determining Light Directions

- Assume orthographic projection
- Viewing direction R = [0,0,-1]
- Normal?

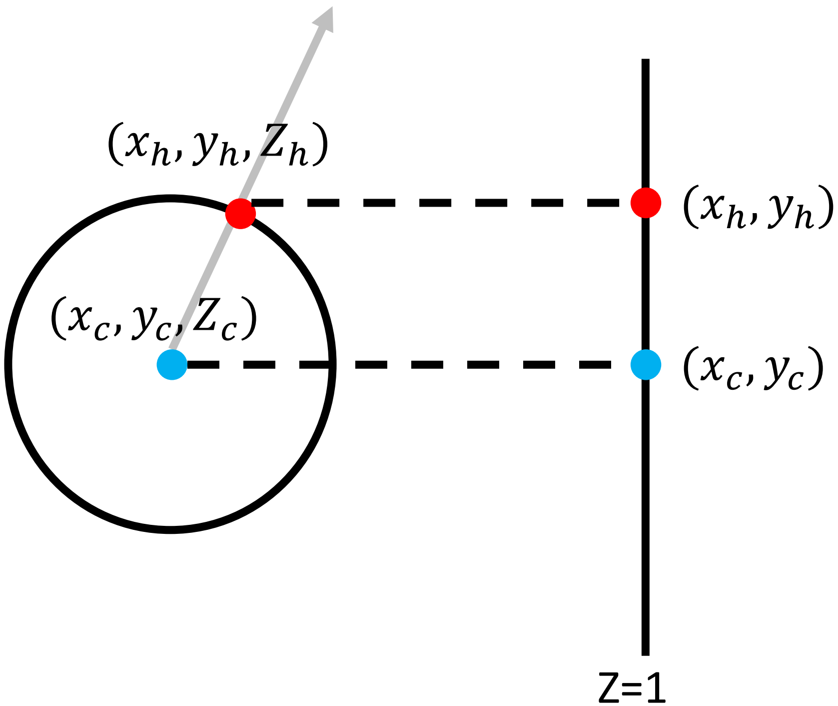

Z_h and Z_c are unknown, but:

$$(x_h - x_c)^2 + (y_h - y_c)^2 + (Z_h - Z_c)^2 = r^2$$

$(Z_h - Z_c)$ can be computed

$(x_h - x_c, y_h - y_c, Z_h - Z_c)$ is the normal

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

Photometric Stereo

What results can you get?

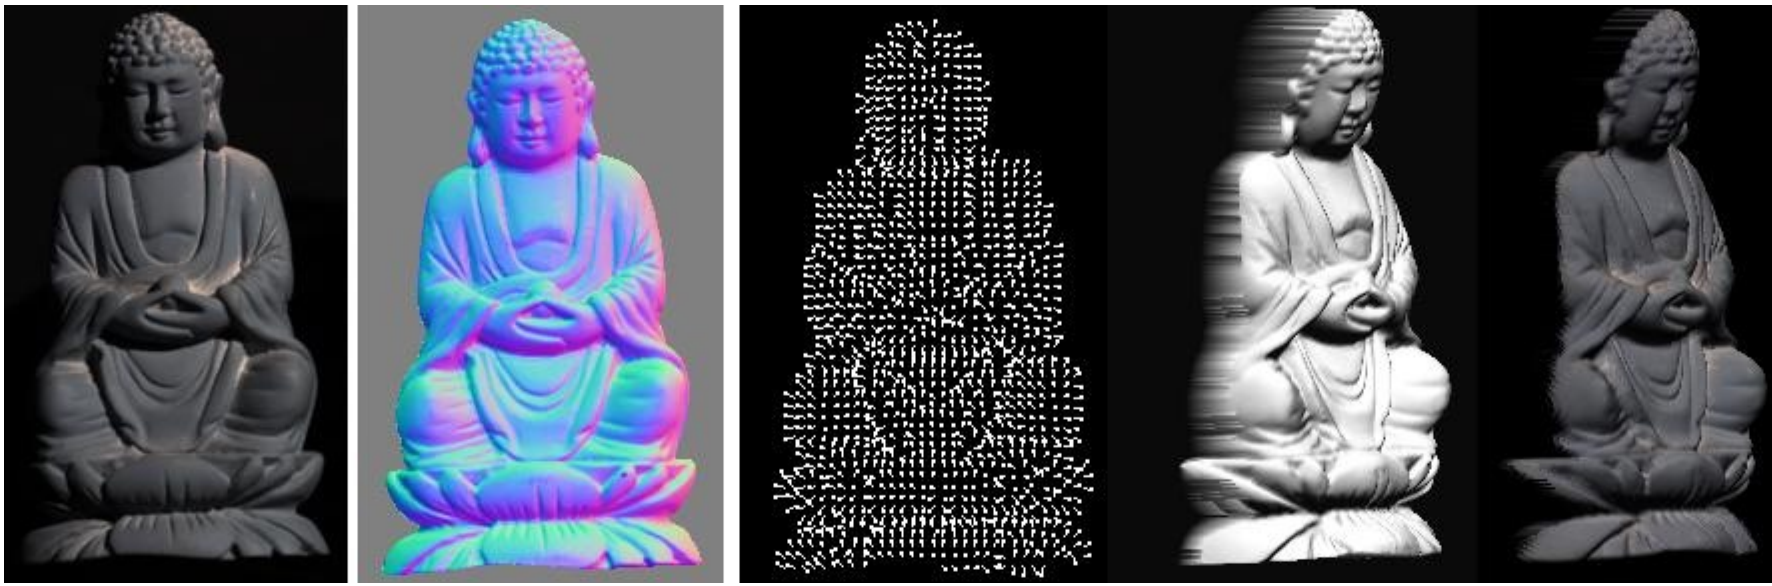

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Results

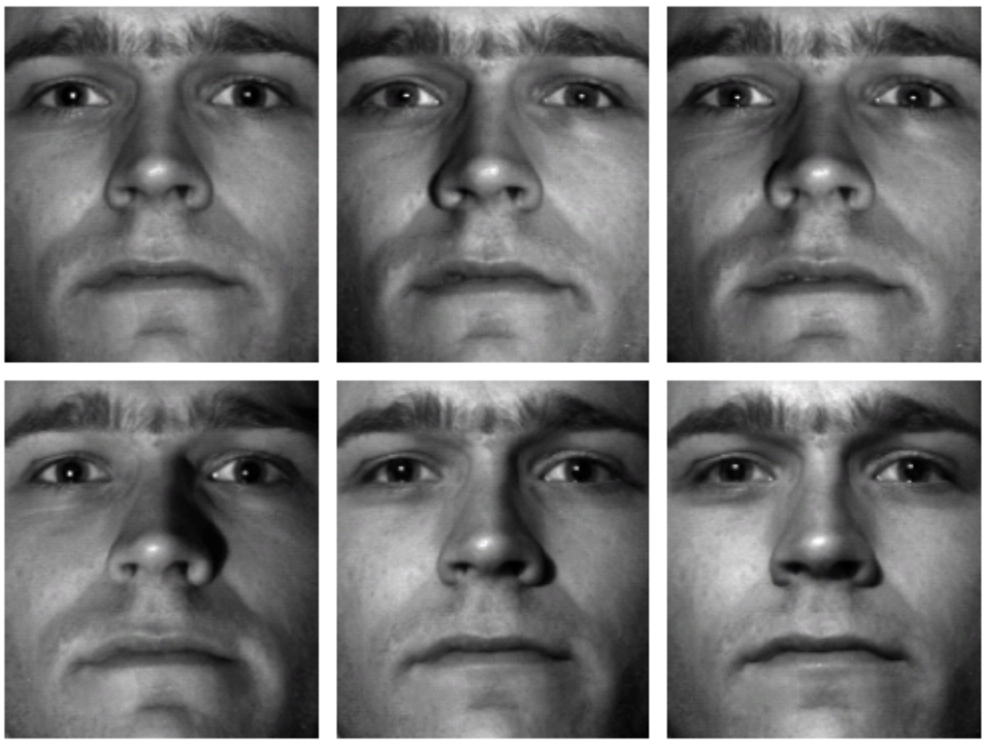

from Athos Georghiades

Results

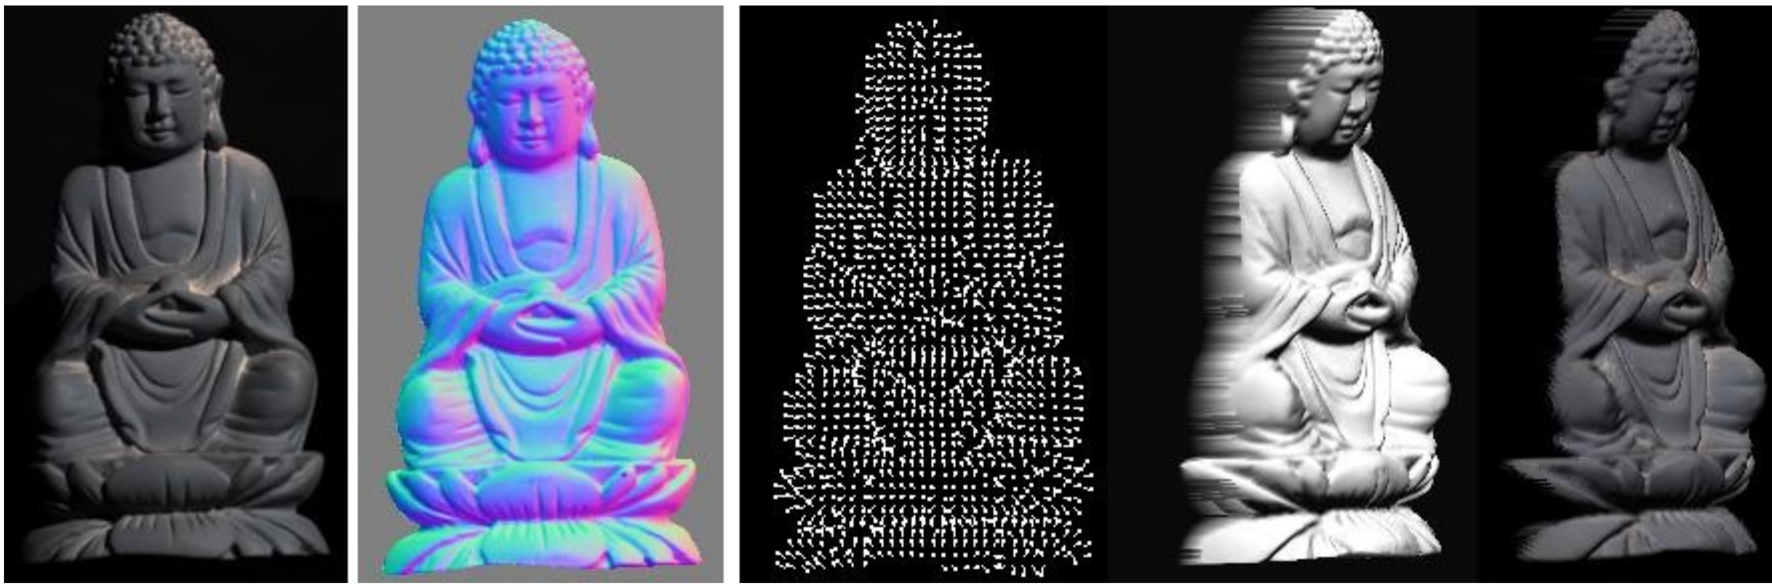

Input (1 of 12) | Normals (RGB colormap) | Normals (vectors) | Shaded 3D rendering | Textured 3D rendering

Photometric stereo

- How many lights / images do you need?
- Are there any constraints on where to put the lights?

Color Images

- Now we have 3 sets of equations for a pixel:

$$\mathbf{I}_r = \rho_r \mathbf{L}_r^T \mathbf{N}$$

$$\mathbf{I}_g = \rho_g \mathbf{L}_g^T \mathbf{N}$$

$$\mathbf{I}_b = \rho_b \mathbf{L}_b^T \mathbf{N}$$

Red values from all images

Green values for all pixels

Blue values from all images

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

Color Images

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

$$Q(\rho) = \sum_i (I_i - \rho \mathbf{L}_i^T \mathbf{N})^2$$

- Want to minimize $Q(\rho)$
  - Take derivative and set to 0

$$\frac{dQ(\rho)}{d\rho} = -2 \sum_i (I_i - \rho \mathbf{L}_i^T \mathbf{N}) \mathbf{L}_i^T \mathbf{N}$$

$$\frac{dQ(\rho)}{d\rho} = 0 \Rightarrow \rho = \frac{\sum_i I_i \mathbf{L}_i^T \mathbf{N}}{\sum_i (\mathbf{L}_i^T \mathbf{N})^2}$$

Question

- How many color images do you need?

- How many color images do you need if the object was grayscale?

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

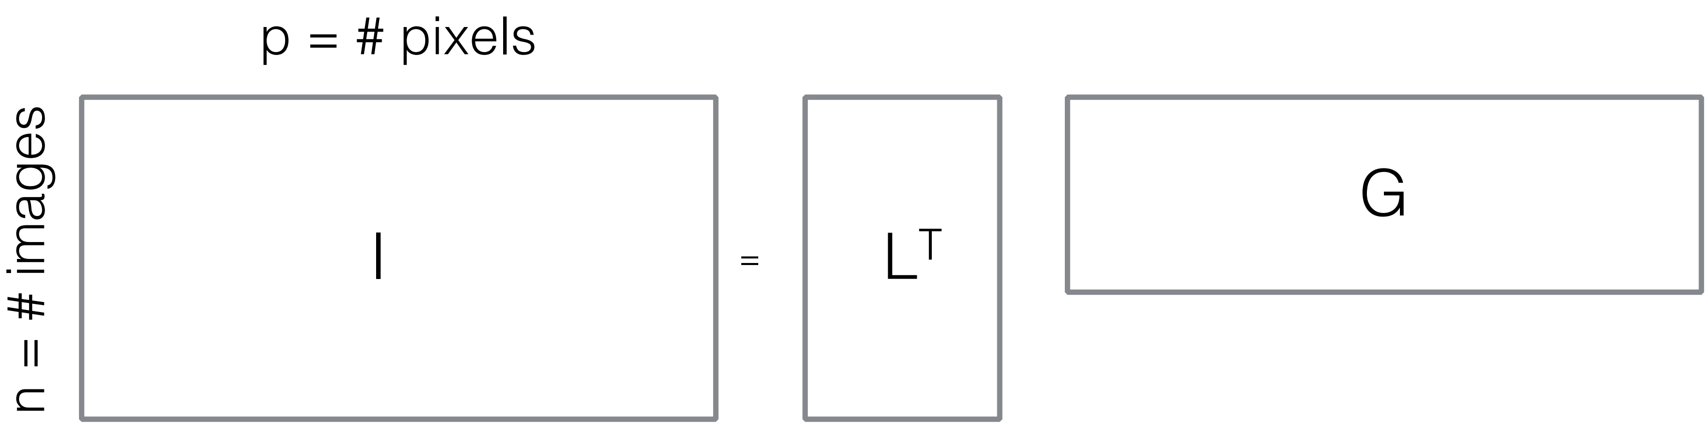

Unknown Lighting

Measurements (one image per row) | Light directions (scaled by intensity) | Surface normals (scaled by albedo)

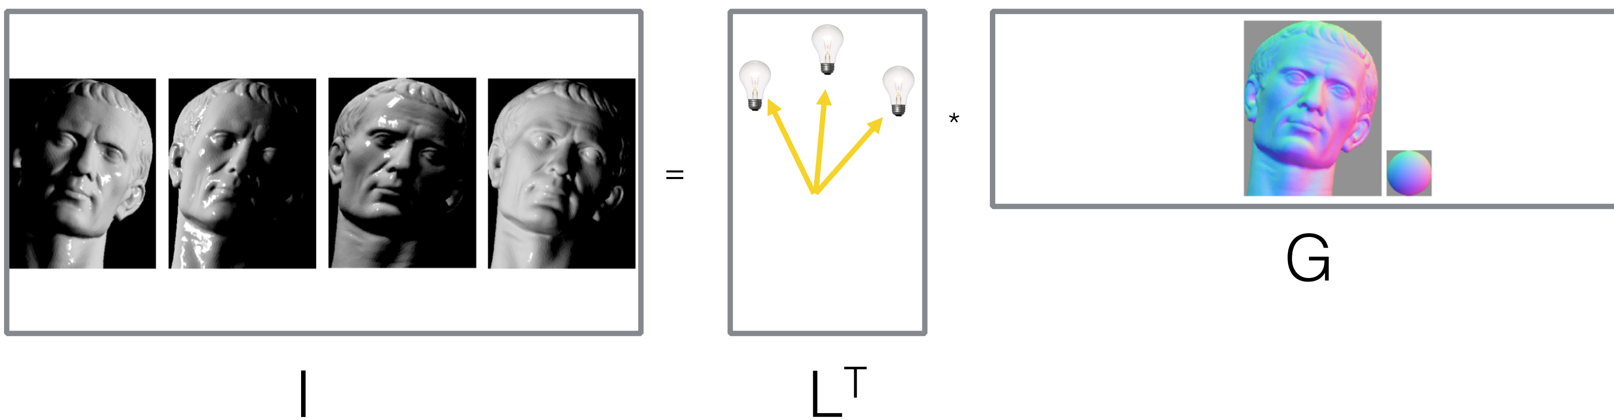

Both L and G are now unknown!

This is a matrix factorization problem.

Unknown Lighting

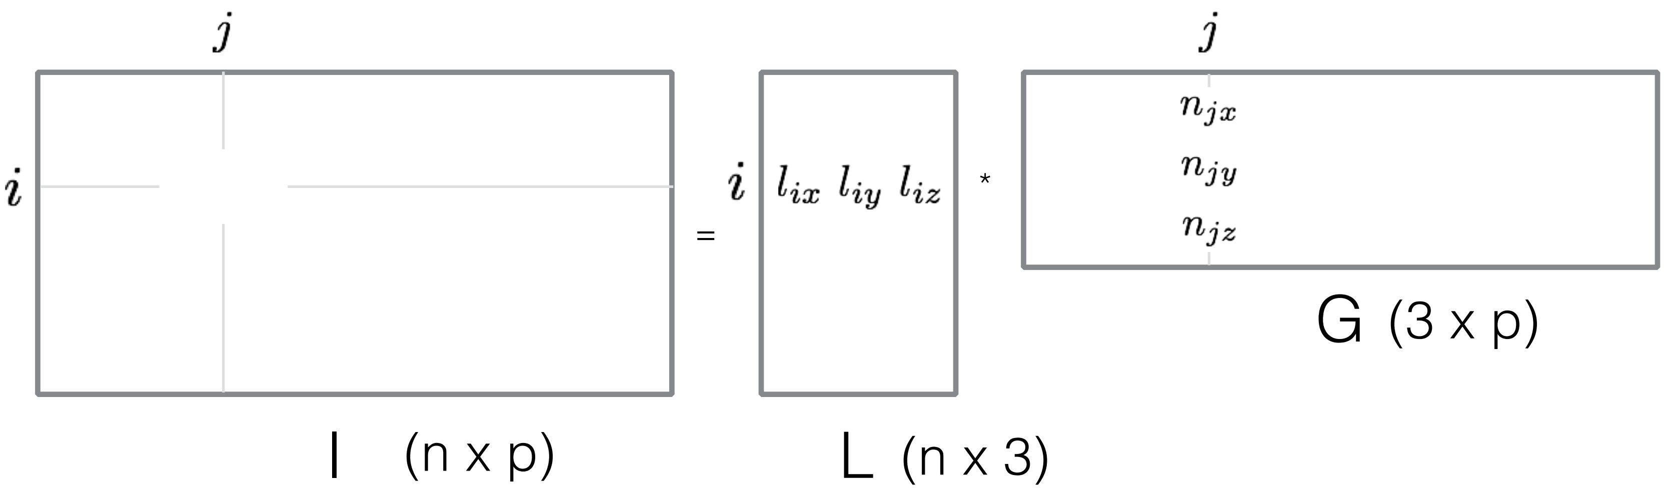

There's hope: We know that I is rank 3

Unknown Lighting

Use the SVD to decompose I:

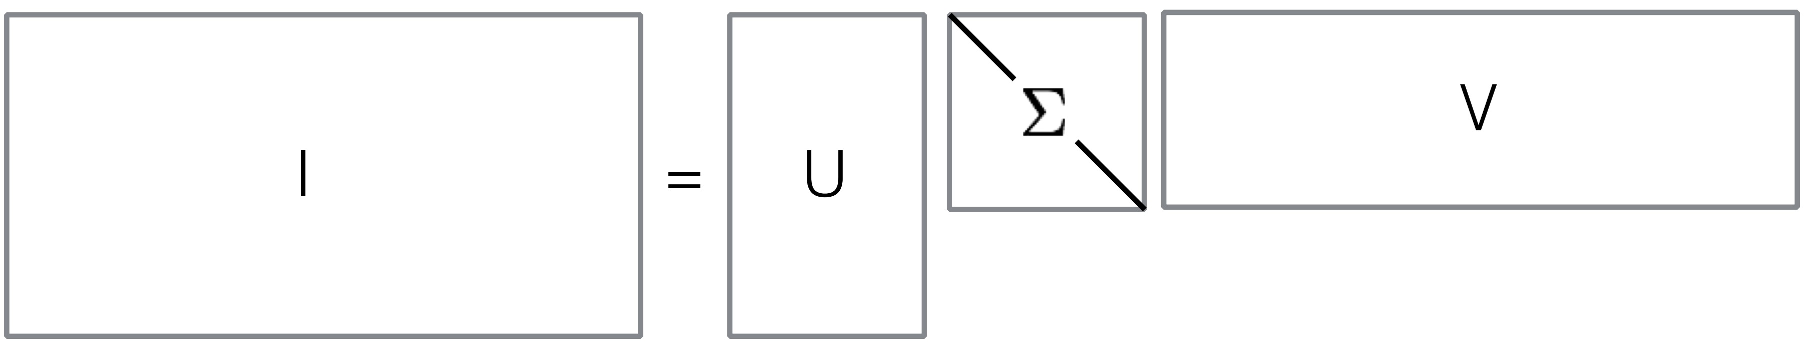

SVD gives the best rank-3 approximation of a matrix.

Unknown Lighting

Use the SVD to decompose I:

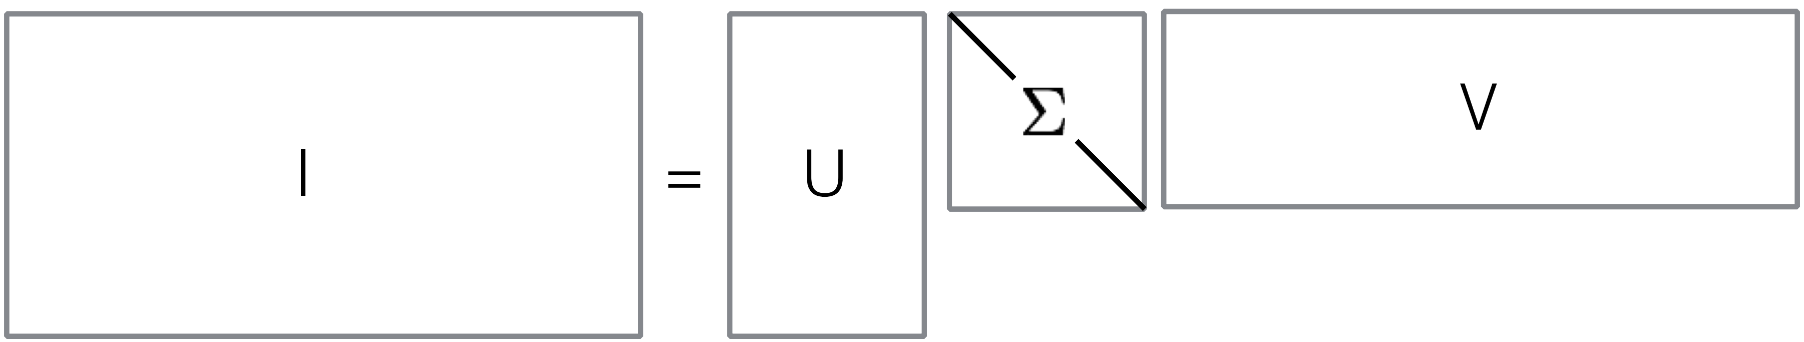

SVD gives the best rank-3 approximation of a matrix.

What do we do with Σ?

Unknown Lighting

Use the SVD to decompose I:

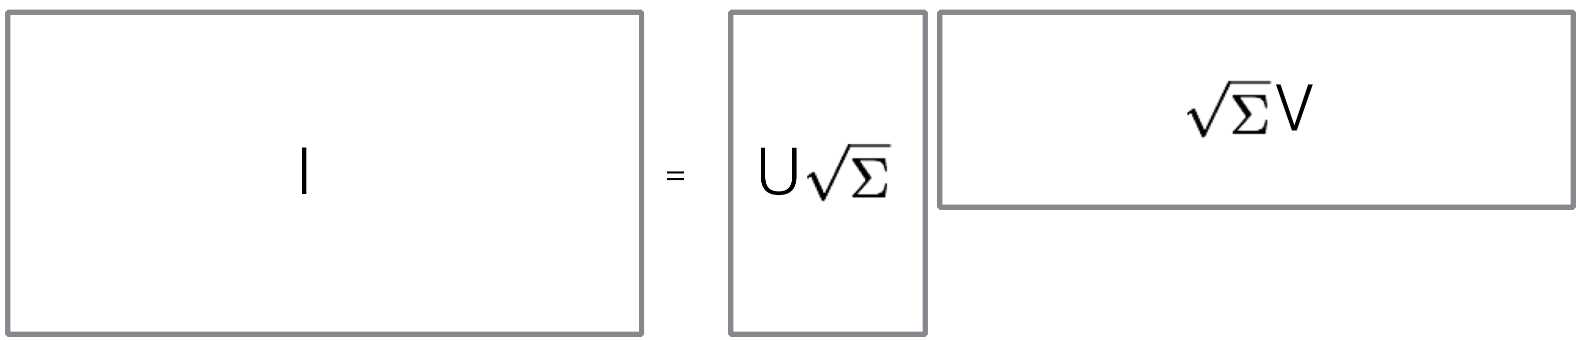

Can we just do that?

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:

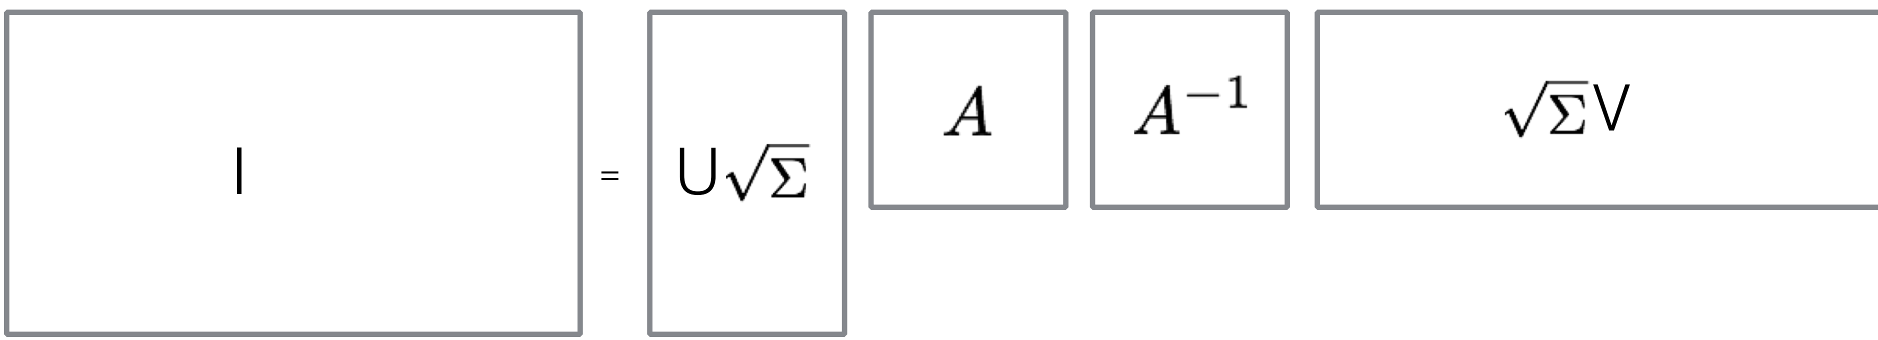

Can we just do that? …almost. $L = U\sqrt{\Sigma}A, G = A^{-1}\sqrt{\Sigma}V$

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

Unknown Lighting

Use the SVD to decompose I:

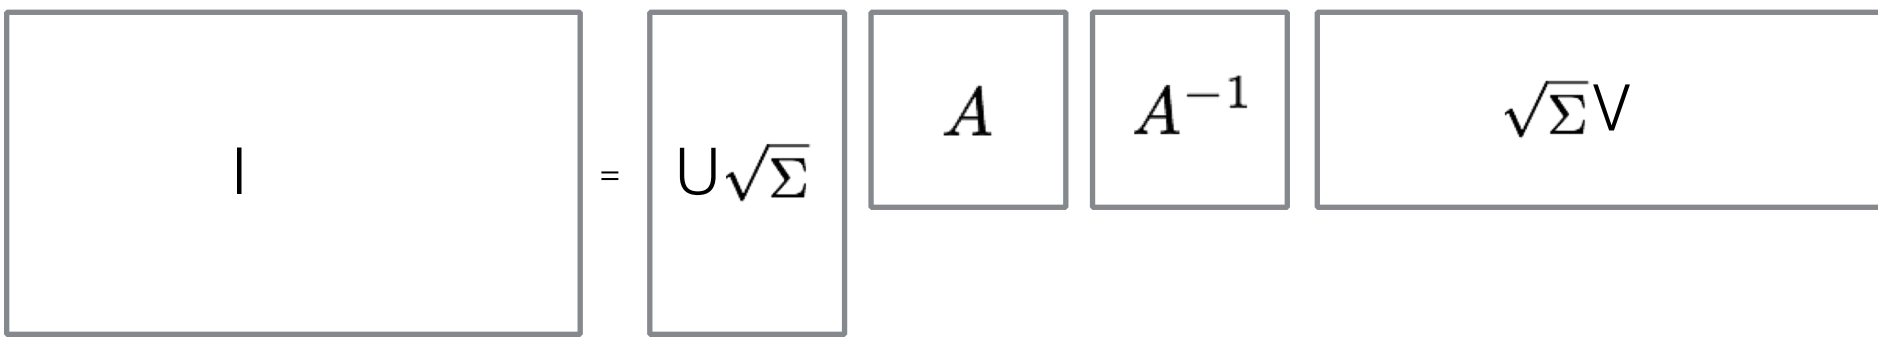

$$L = U\sqrt{\Sigma}A, G = A^{-1}\sqrt{\Sigma}V$$

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.

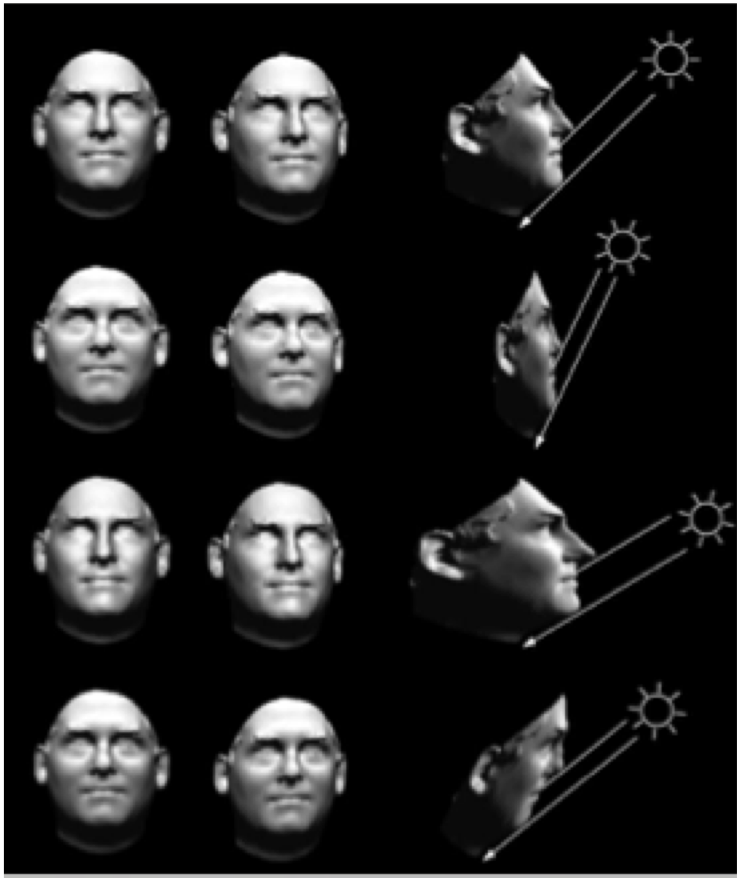

[Belhumeur et al.'97]

Photometric stereo

- How do we deal with non-lambertian surfaces?
- How do we make sure we know the lighting directions?
- How can we effectively use color?

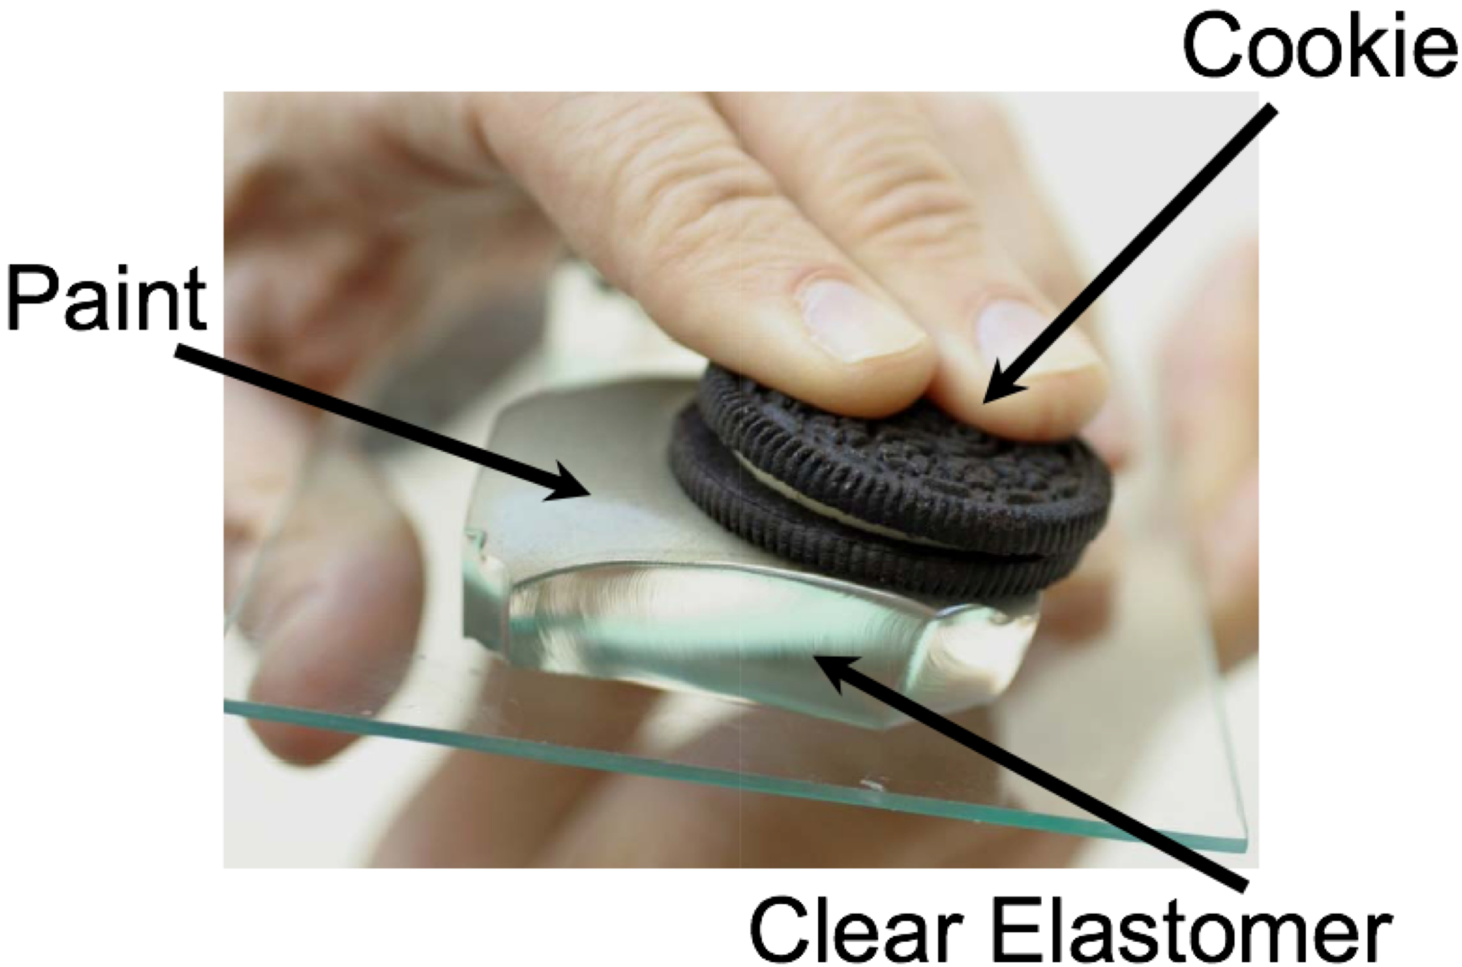

Johnson and Adelson, 2009


OREO


OREO

Lights, camera, action

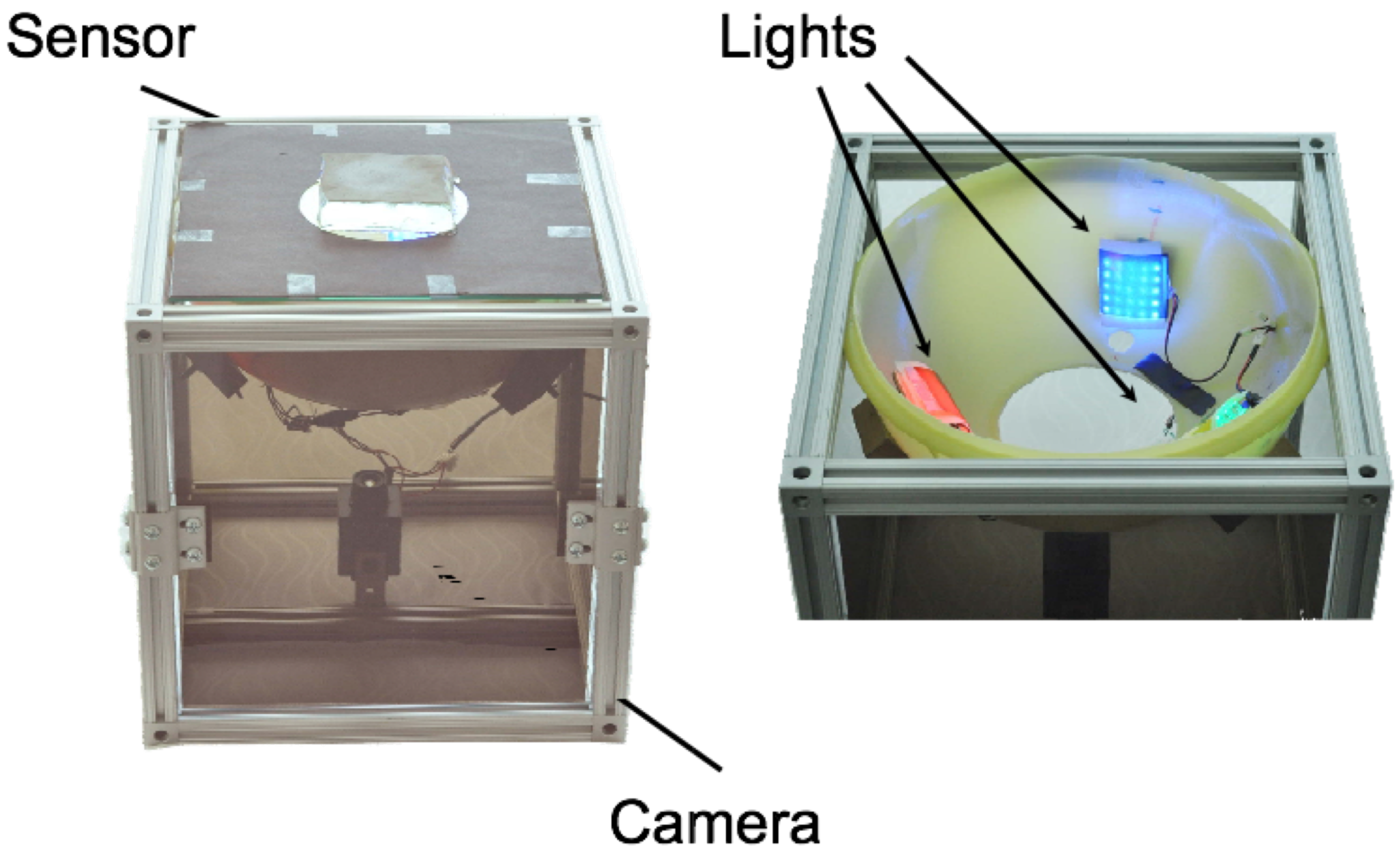

(a) (b)

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.

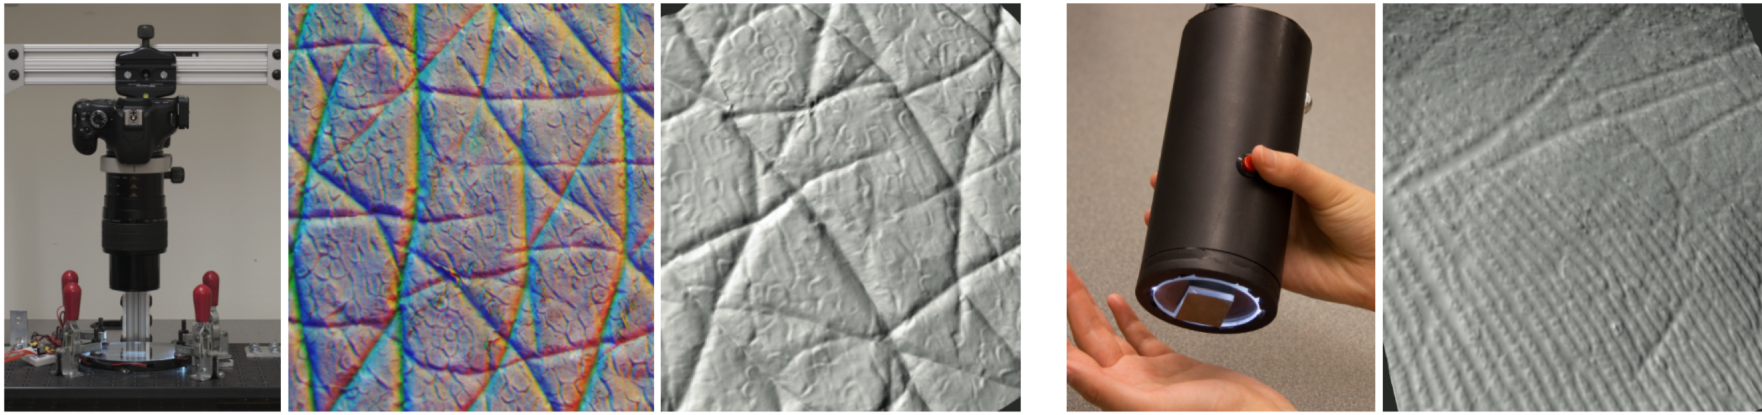

(a) bench configuration (b) captured (c) reconstruction (d) portable configuration (e) reconstruction

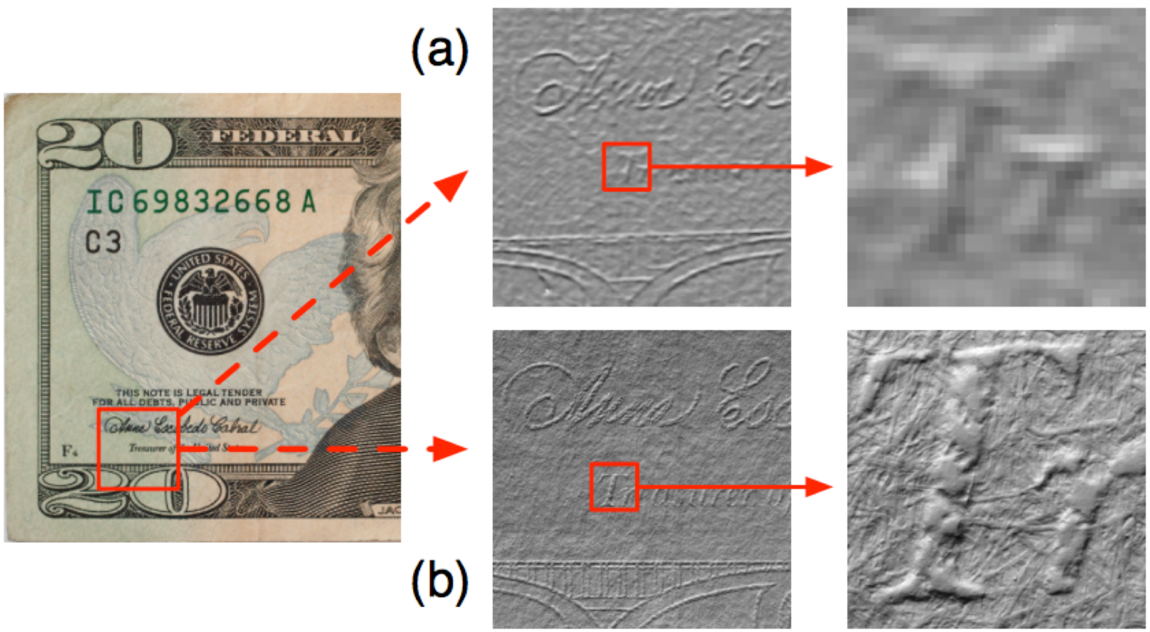

Figure 7: *Comparison with the high-resolution result from the original retrographic sensor. (a) Rendering of the high-resolution \$20 bill example from the original retrographic sensor with a close-up view. (b) Rendering of the captured geometry using our method.*

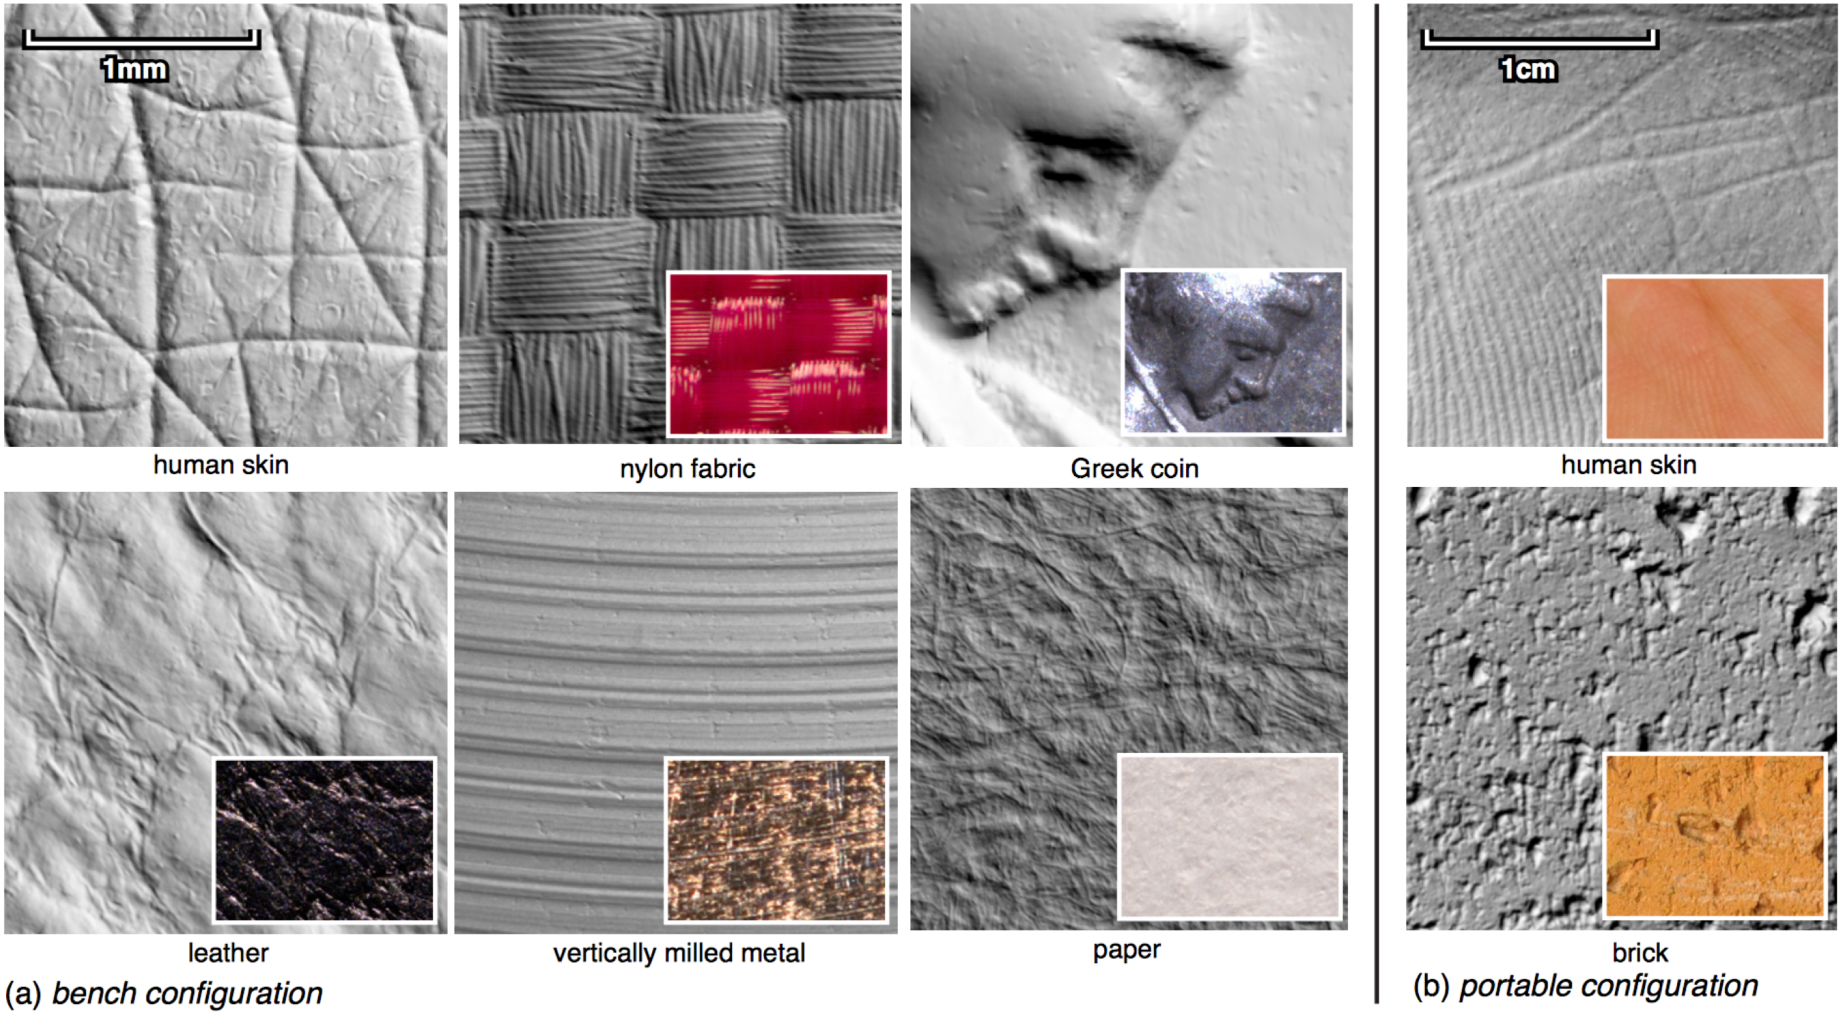

Figure 9: *Example geometry measured with the bench and portable configurations. Outer image: rendering under direct lighting. Inset: macro photograph of original sample. Scale shown in upper left. Color images are shown for context and are to similar, but not exact scale.*

Summary of reconstruction

- Given multiple cameras, correspondence, cameras
  - Triangulation
  - Set up linear equations of the form $Ax = 0$
  - Correspondence must satisfy epipolar constraint
- Stereo is easy for rectified stereo cameras
  - Only disparities along a row
  - Disparity direct measure of depth
- Photometric stereo: use different light sources
  - Gives you normals and albedo