

CS5670 : Computer Vision

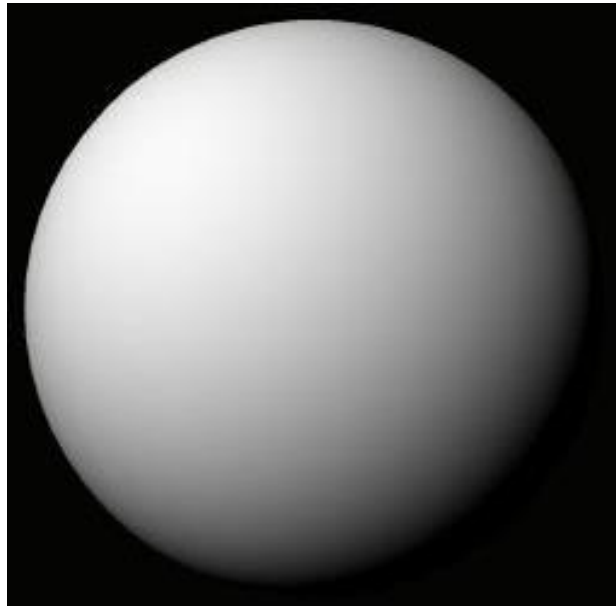
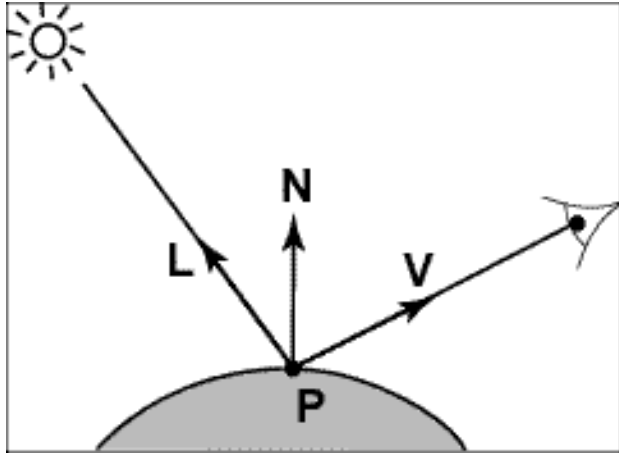
Photometric stereo



Announcements

- Project 4 (stereo) released today, due Friday, April 1, at 8pm
- To be done in groups of 2

Recap: Lambertian (Diffuse) Reflectance

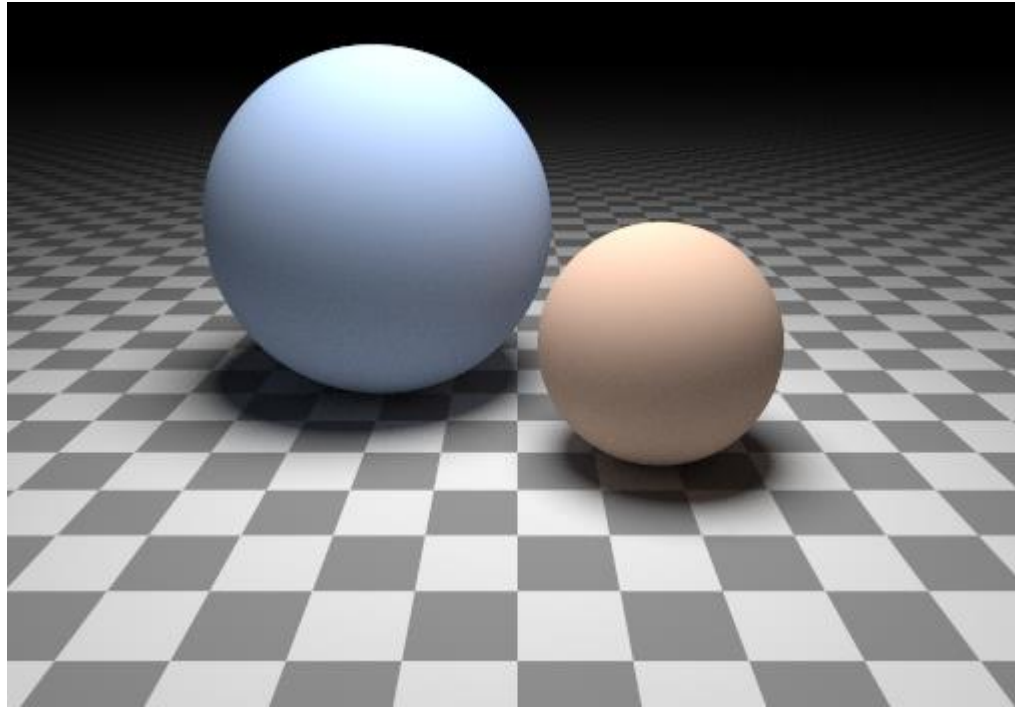


$$I = k_d \mathbf{N} \cdot \mathbf{L}$$

- I : observed image intensity
- k_d : object albedo
- \mathbf{N} : surface normal
- \mathbf{L} : light source direction

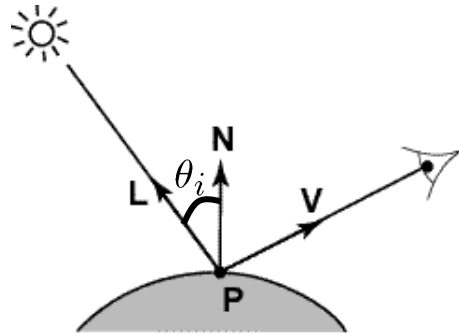
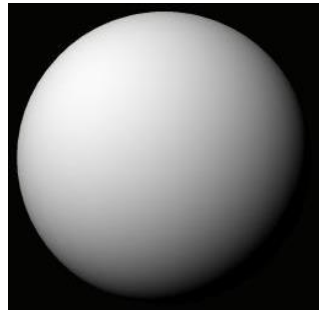
← Lambertian sphere with constant albedo lit by a directional light source

Can we determine shape from lighting?



- Are these spheres?
 - Or just flat discs painted with varying albedo?

A Single Image: Shape from shading



Suppose (for now) $k_d = 1$

$$\begin{aligned} I &= k_d \mathbf{N} \cdot \mathbf{L} \\ &= \mathbf{N} \cdot \mathbf{L} \\ &= \cos \theta_i \end{aligned}$$

You can directly measure angle between normal and light source

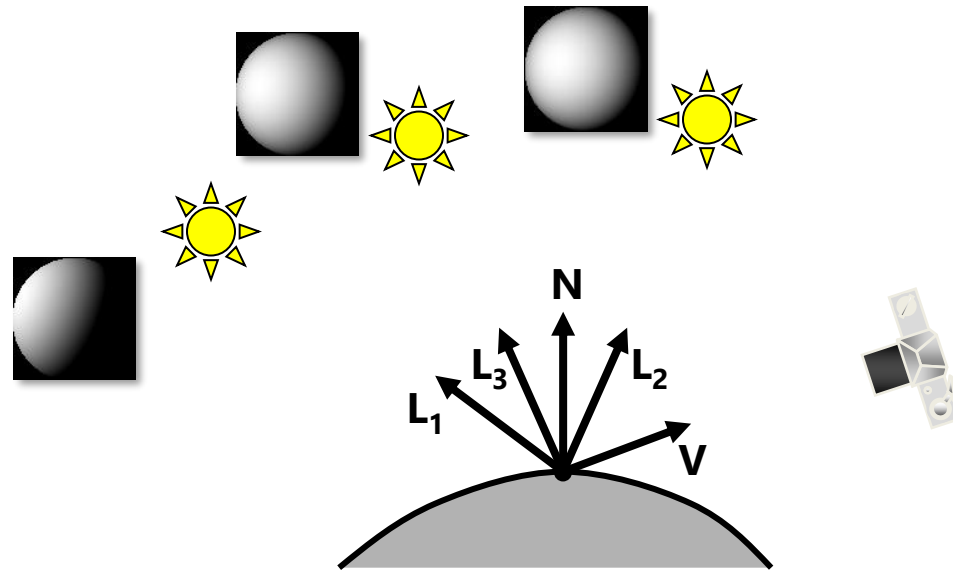
- Not quite enough information to compute surface shape
- But can be if you add some additional info, for example
 - assume a few of the normals are known (e.g., along silhouette)
 - constraints on neighboring normals—“integrability”
 - smoothness
- Hard to get it to work well in practice
 - plus, how many real objects have constant albedo?
 - But, deep learning can help



<https://www.good.is/optical-illusion-plates-and-bowls-upside-down-or-not>

Let's take more than one photo!

Photometric stereo



$$I_1 = k_d \mathbf{N} \cdot \mathbf{L}_1$$

$$I_2 = k_d \mathbf{N} \cdot \mathbf{L}_2$$

$$I_3 = k_d \mathbf{N} \cdot \mathbf{L}_3$$

Can write this as a matrix equation:

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = k_d \begin{bmatrix} \mathbf{L}_1^T \\ \mathbf{L}_2^T \\ \mathbf{L}_3^T \end{bmatrix} \mathbf{N}$$

Solving the equations

$$\underbrace{\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix}}_{\mathbf{I} \quad 3 \times 1} = \underbrace{\begin{bmatrix} \mathbf{L}_1^T \\ \mathbf{L}_2^T \\ \mathbf{L}_3^T \end{bmatrix}}_{\mathbf{L} \quad 3 \times 3} \underbrace{k_d \mathbf{N}}_{\mathbf{G} \quad 3 \times 1}$$
$$\mathbf{G} = \mathbf{L}^{-1} \mathbf{I}$$
$$k_d = \|\mathbf{G}\|$$
$$\mathbf{N} = \frac{1}{k_d} \mathbf{G}$$

Solve one such linear system **per pixel** to solve for that pixel's surface normal

More than three lights

Can get better results by using more than 3 lights

$$\begin{bmatrix} I_1 \\ \vdots \\ I_n \end{bmatrix} = \begin{bmatrix} \mathbf{L}_1 \\ \vdots \\ \mathbf{L}_n \end{bmatrix} k_d \mathbf{N}$$

Least squares solution:

$$\begin{aligned} \mathbf{I} &= \mathbf{L}\mathbf{G} \\ \mathbf{L}^T \mathbf{I} &= \mathbf{L}^T \mathbf{L}\mathbf{G} \\ \mathbf{G} &= (\mathbf{L}^T \mathbf{L})^{-1} (\mathbf{L}^T \mathbf{I}) \end{aligned}$$

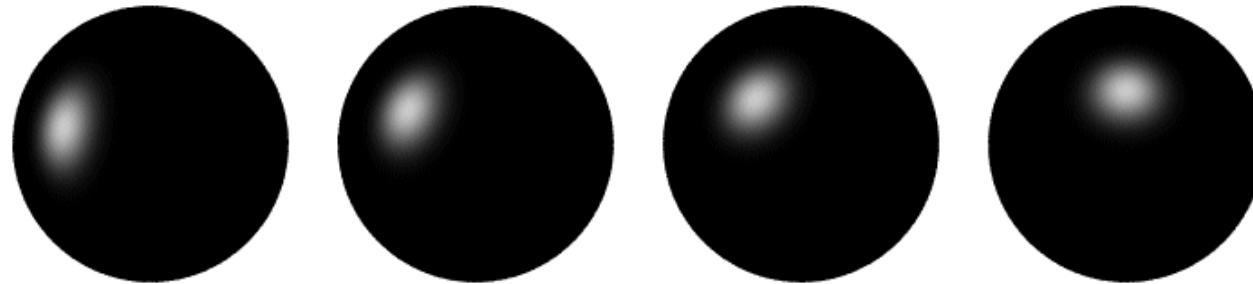
Solve for \mathbf{N} , k_d as before

What's the size of $\mathbf{L}^T \mathbf{L}$?



Computing light source directions

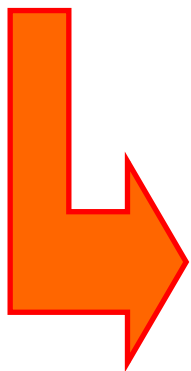
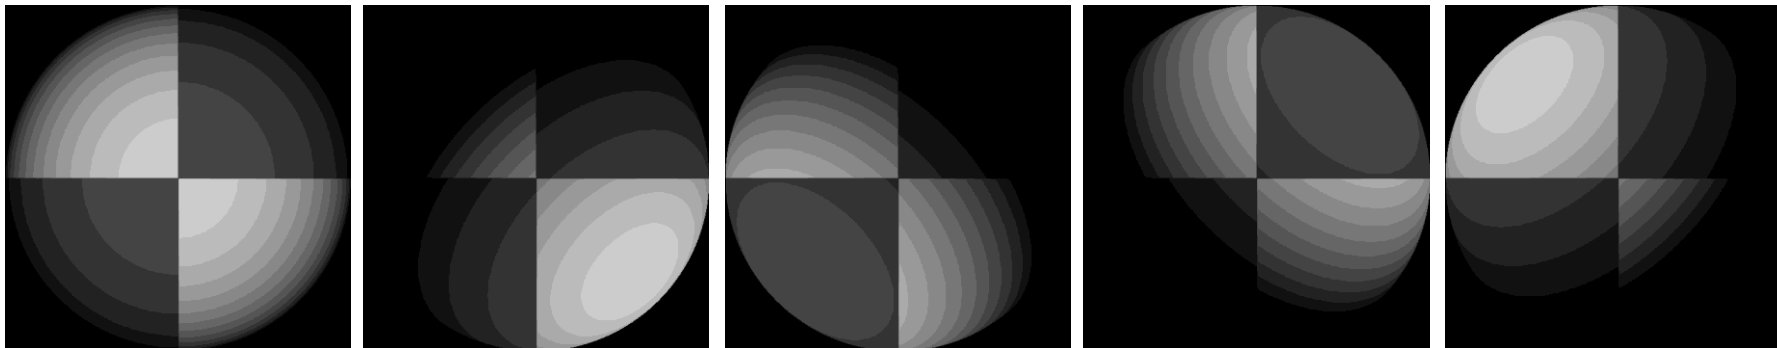
Trick: place a chrome sphere in the scene



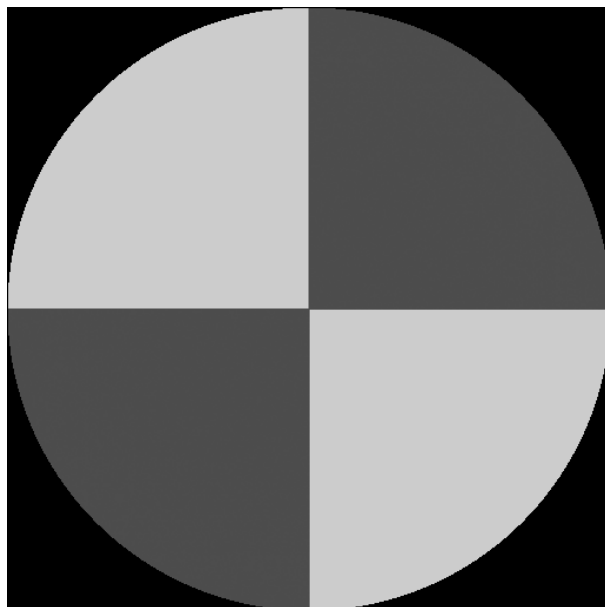
– the location of the highlight tells you where the light source is

Example

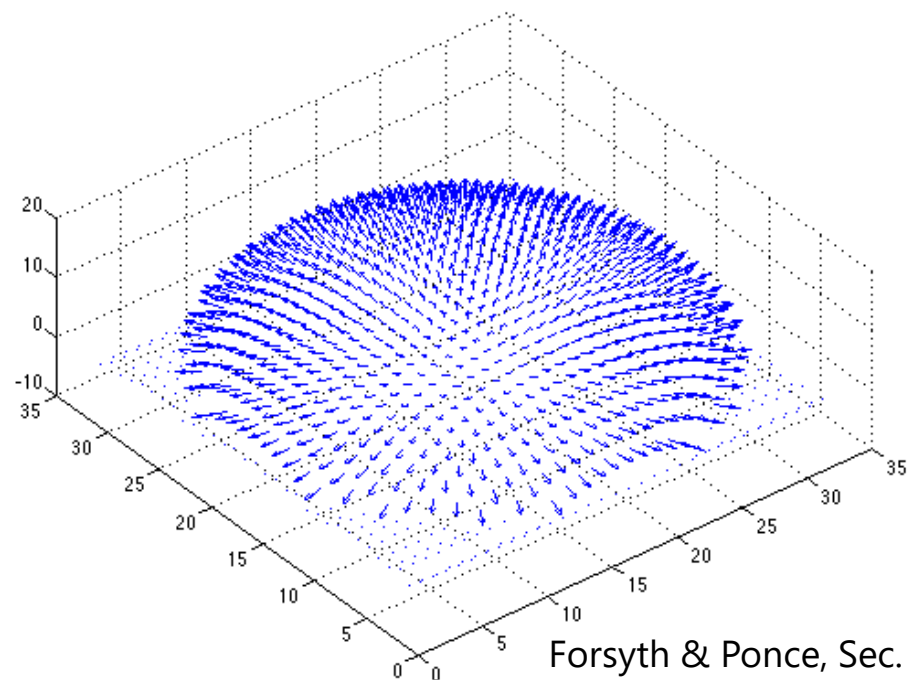
Input views



Recovered albedo



Recovered normal field



Depth from normals

- Solving the linear system per-pixel gives us an estimated surface normal for each pixel
- How can we compute depth from normals?
 - Normals are like the “derivative” of the true depth



Input photo



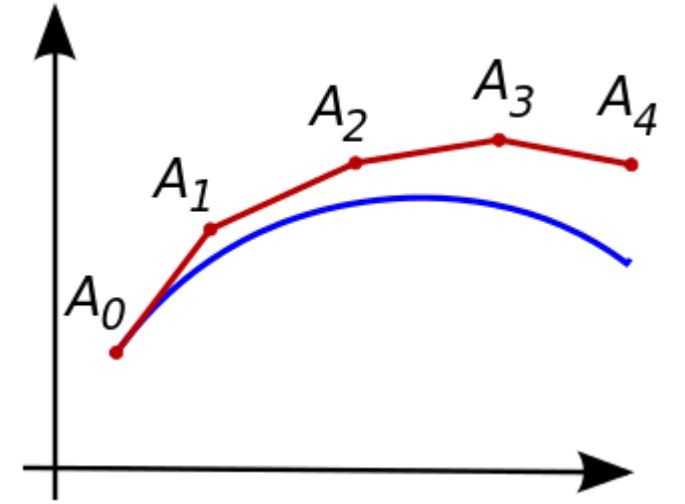
Estimated normals



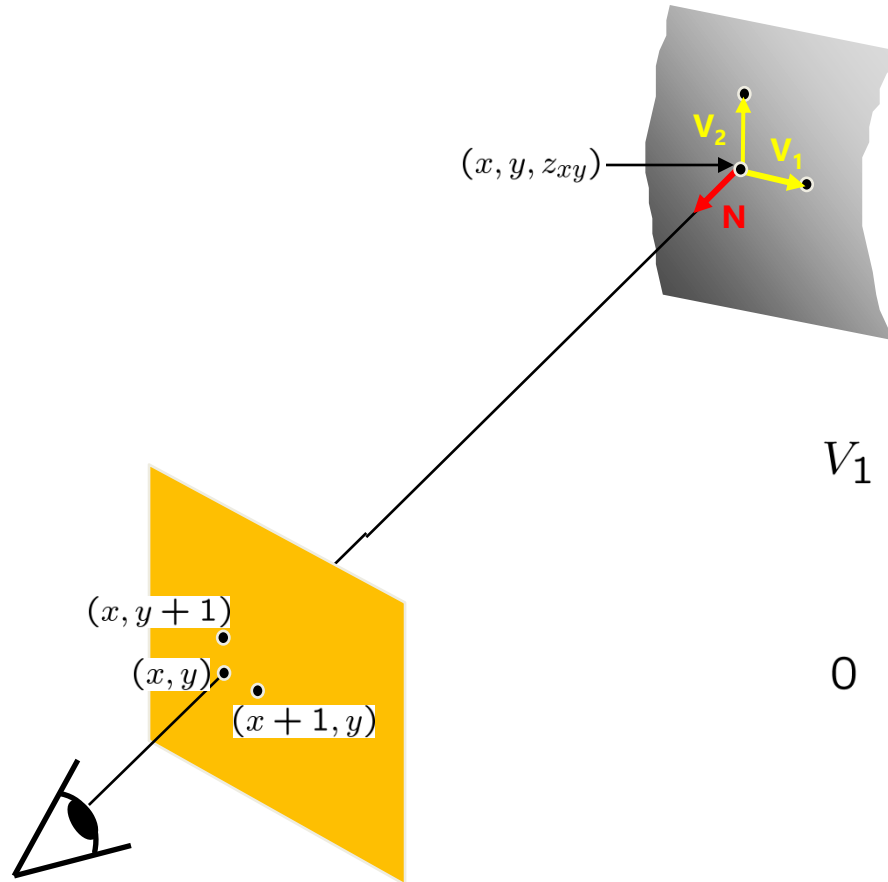
Estimated normals
(needle diagram)

Normal Integration

- Integrating a set of derivatives is easy in 1D
 - (similar to Euler's method from diff. eq. class)
- Could integrate normals in each column / row separately
 - Wouldn't give a good surface
- Instead, we formulate as a linear system and solve for depths that *best agree with the surface normals*



Depth from normals



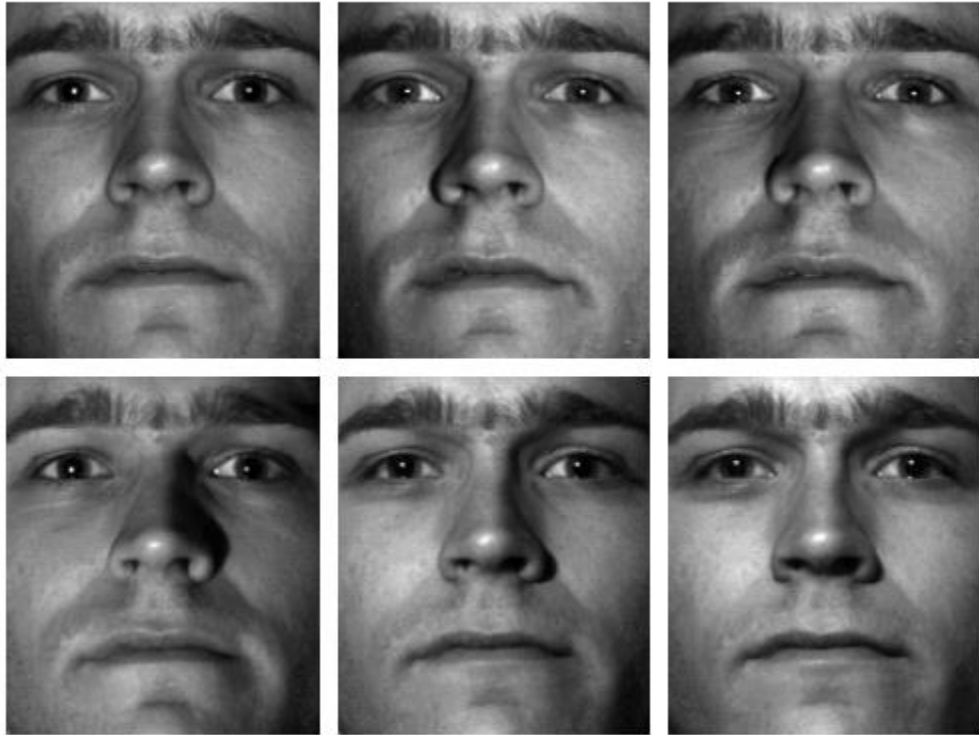
$$\begin{aligned}V_1 &= (x + 1, y, z_{x+1,y}) - (x, y, z_{xy}) \\ &= (1, 0, z_{x+1,y} - z_{xy})\end{aligned}$$

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

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

Results



from Athos Georghiades

Results



Extension

- Photometric Stereo from Colored Lighting

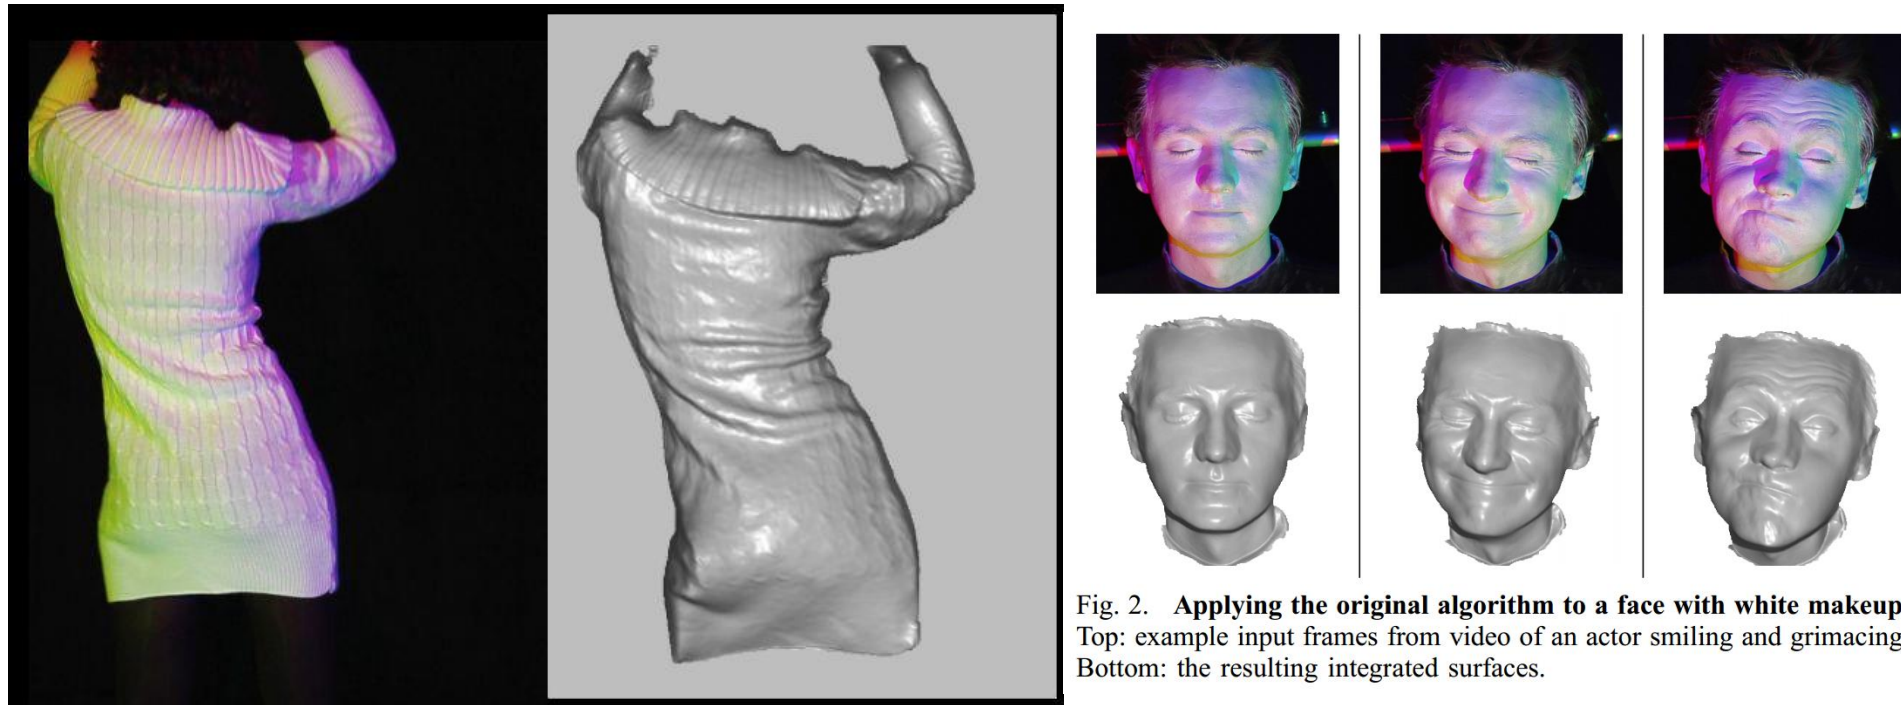


Fig. 2. Applying the original algorithm to a face with white makeup. Top: example input frames from video of an actor smiling and grimacing. Bottom: the resulting integrated surfaces.

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.

Questions?

For now, ignore specular reflection



And Refraction...

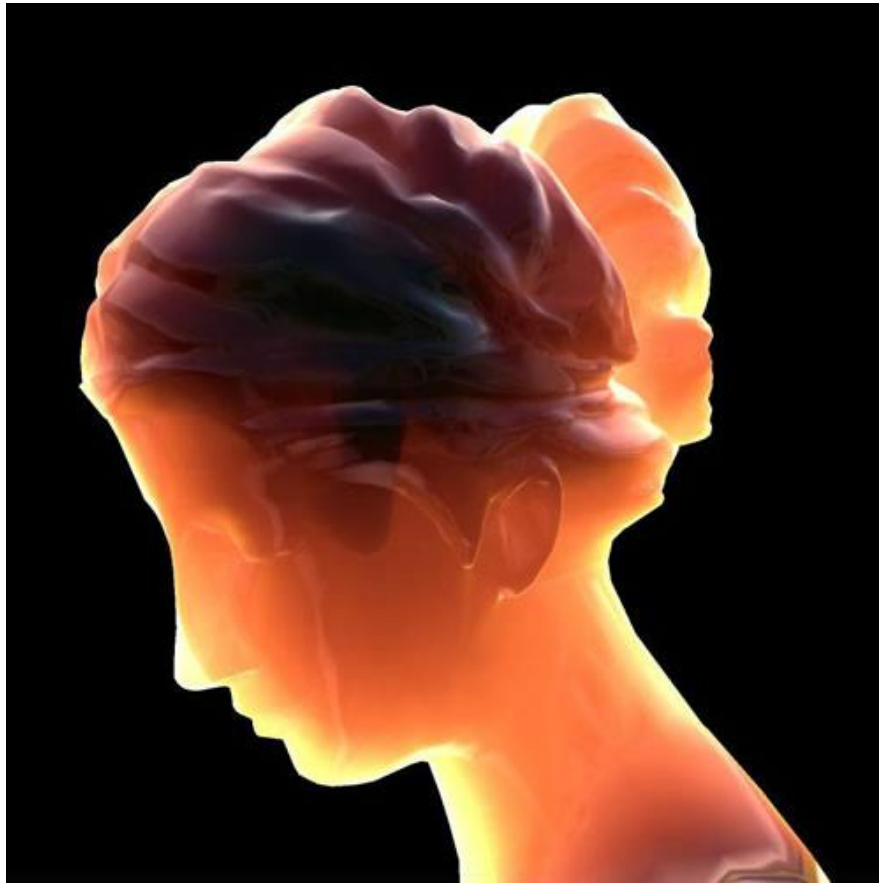


And Interreflections...



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

And Subsurface Scattering...



Limitations

Bigger problems

- doesn't work for shiny things, semi-translucent things
- shadows, inter-reflections

Smaller problems

- camera and lights have to be distant
- calibration requirements
 - measure light source directions, intensities
 - camera response function

Newer work addresses some of these issues

Some pointers for further reading:

- Zickler, Belhumeur, and Kriegman, "[Helmholtz Stereopsis: Exploiting Reciprocity for Surface Reconstruction](#)." IJCV, Vol. 49 No. 2/3, pp 215-227.
- Hertzmann & Seitz, "[Example-Based Photometric Stereo: Shape Reconstruction with General, Varying BRDFs](#)." IEEE Trans. PAMI 2005

GELS*i*GHT

[HOME](#)

[PRODUCTS](#)

[VIDEOS](#)

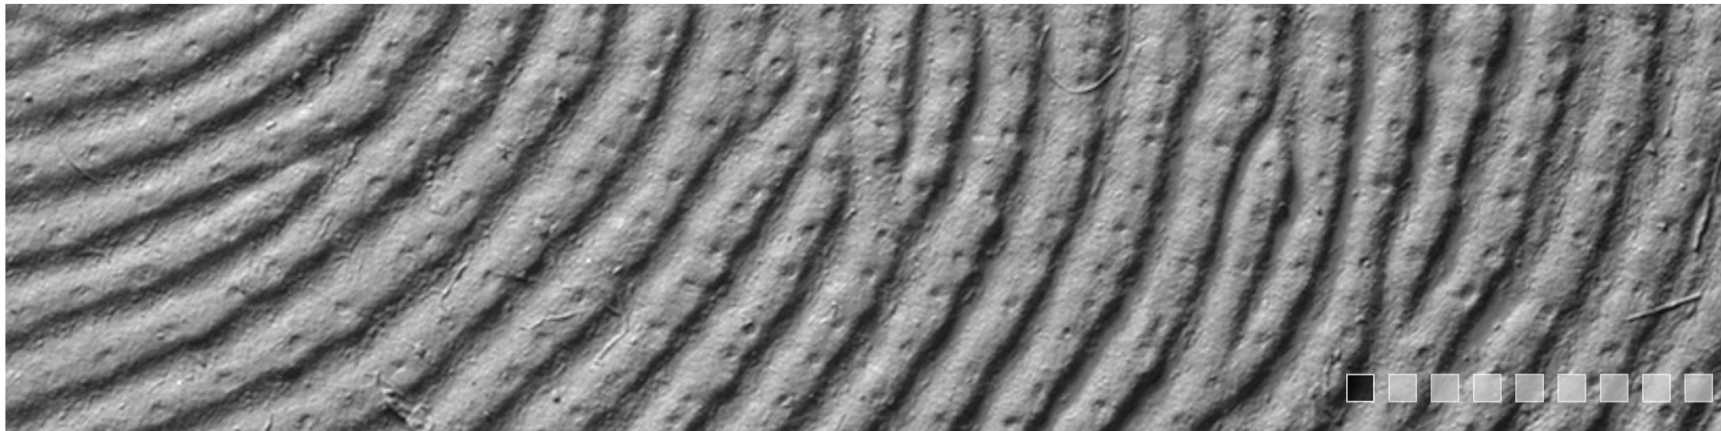
[IMAGES](#)

[PAPERS](#)

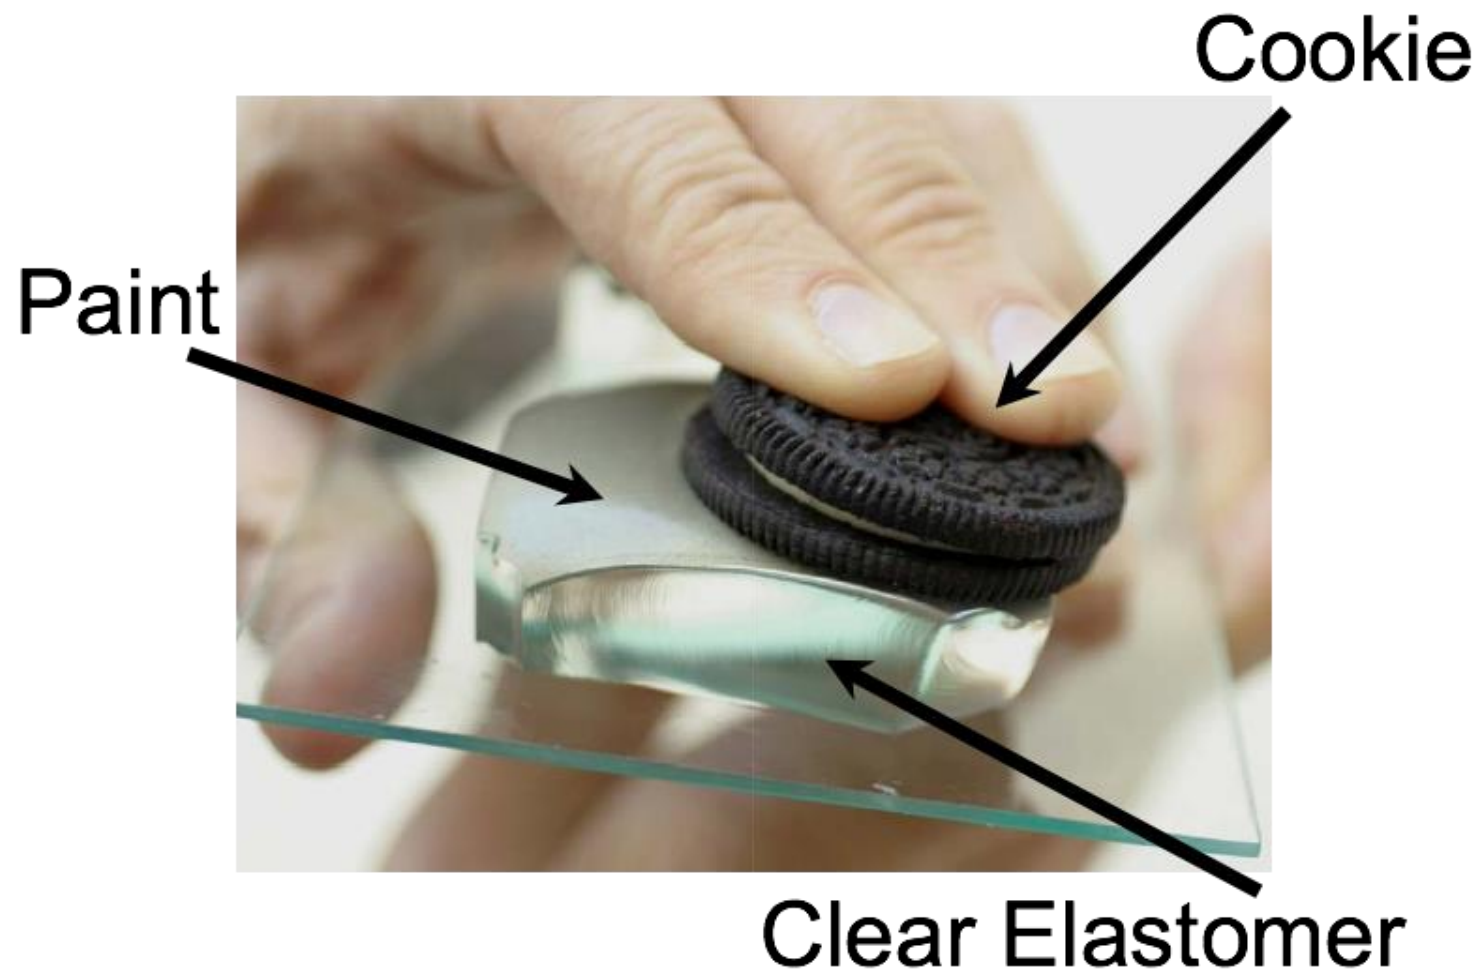
[NEWS](#)

[ABOUT US](#)

[CONTACT](#)



Johnson and Adelson, 2009



Johnson and Adelson, 2009





Lights, camera, action

Sensor



Lights



Camera



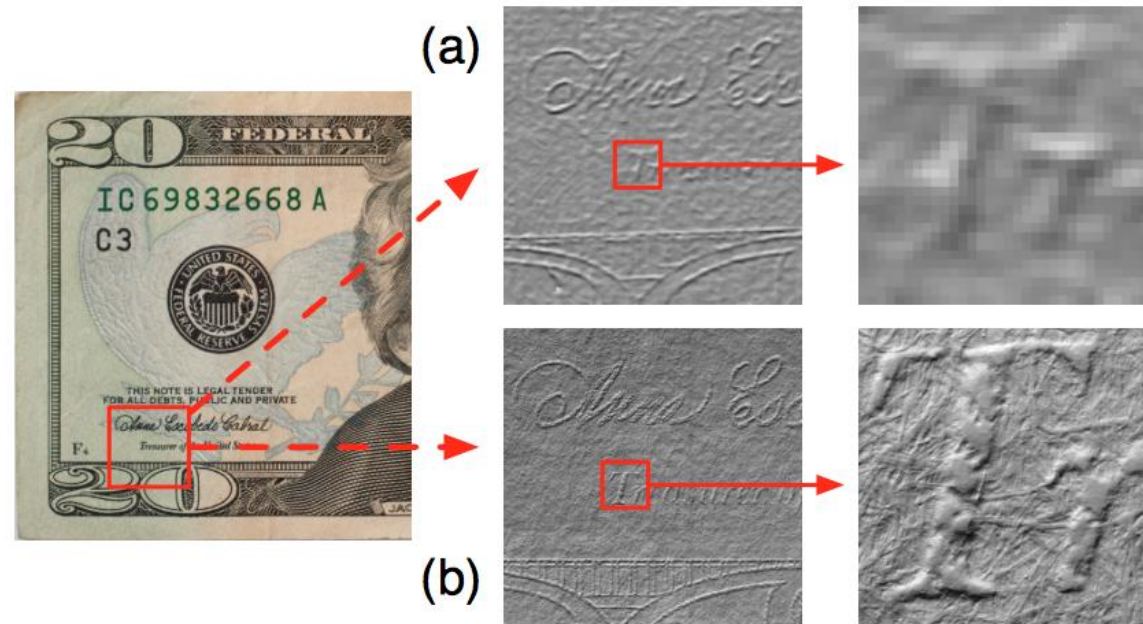
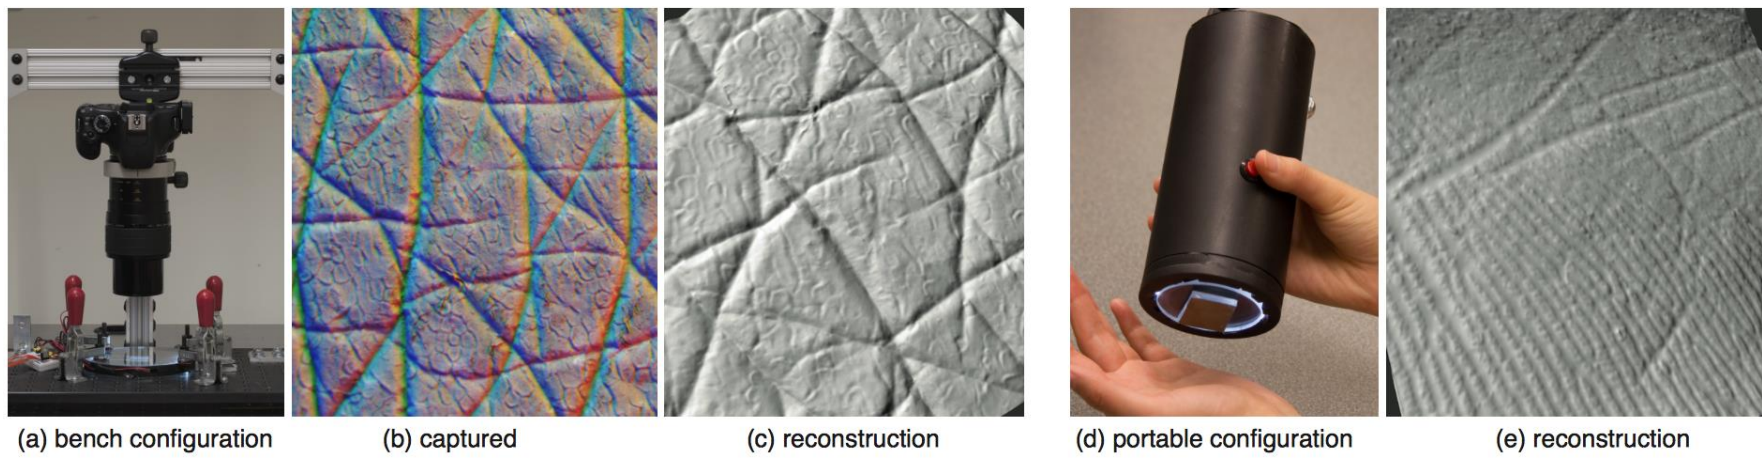


Figure 7: Comparison with the high-resolution result from the original retrographic sensor. (a) Rendering of the high-resolution \$20 bill example from the original retrographic sensor with a close-up view. (b) Rendering of the captured geometry using our method.

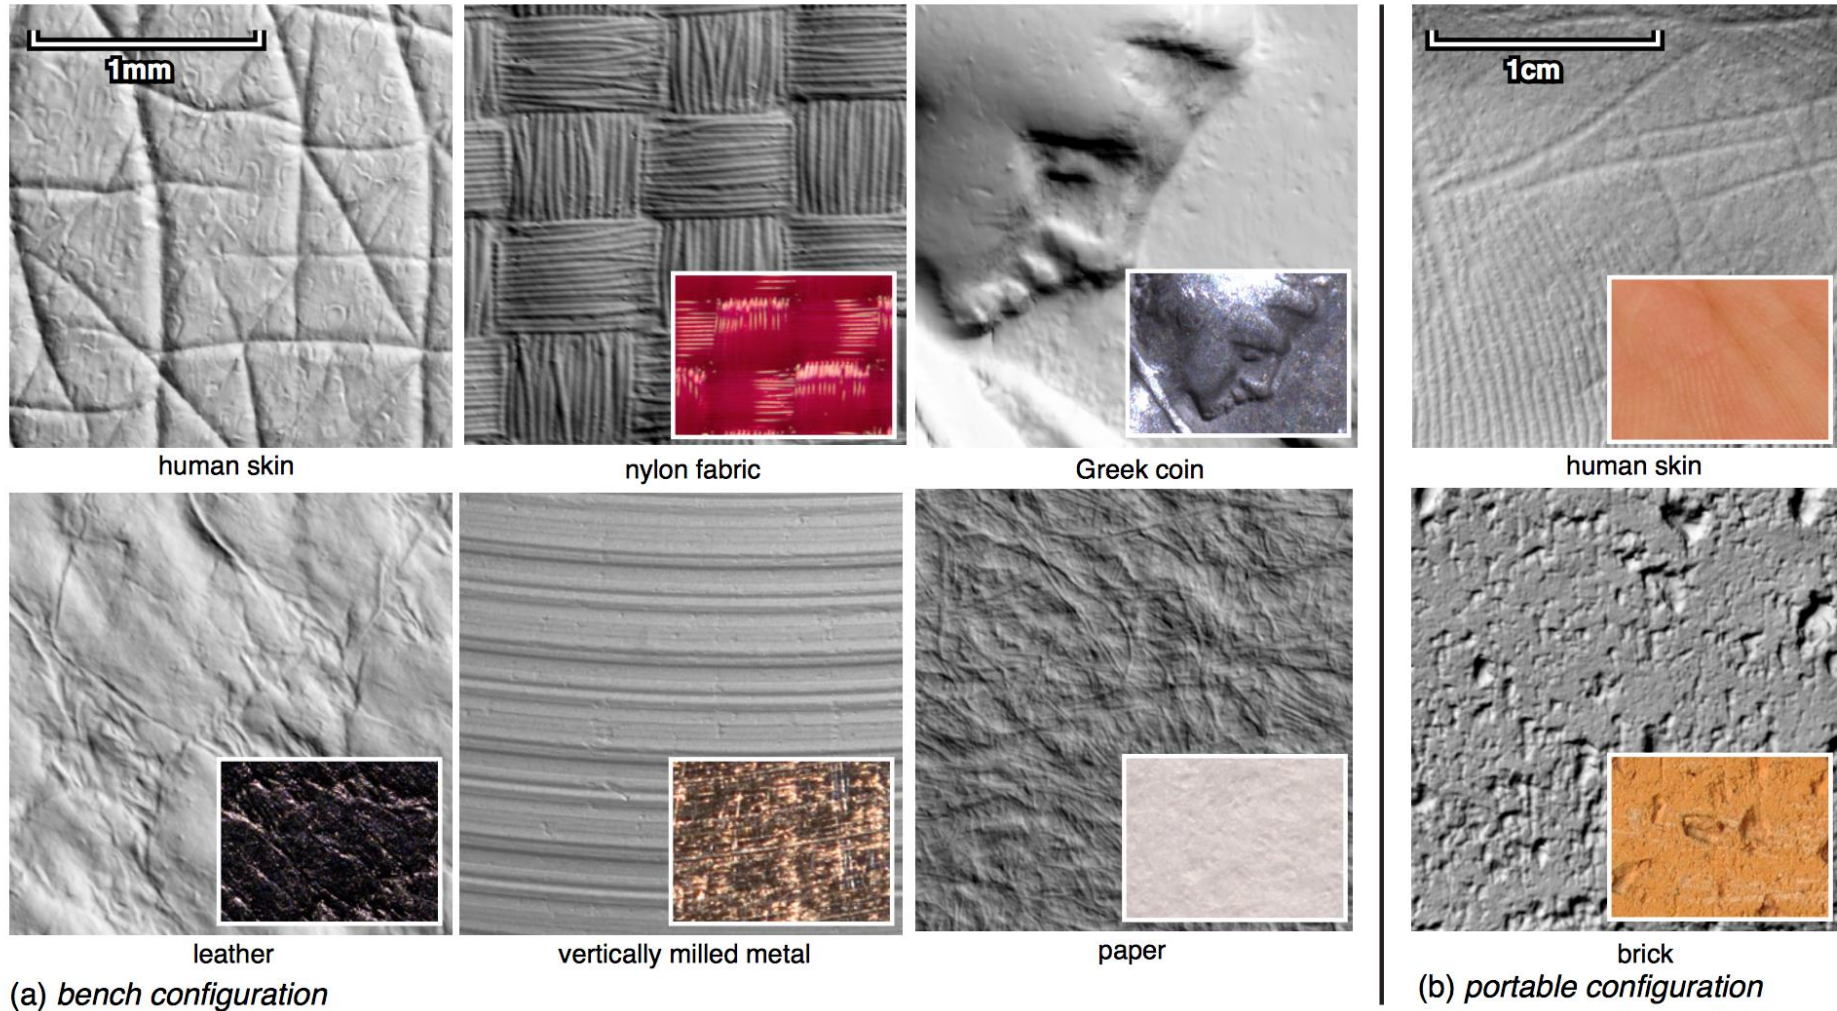
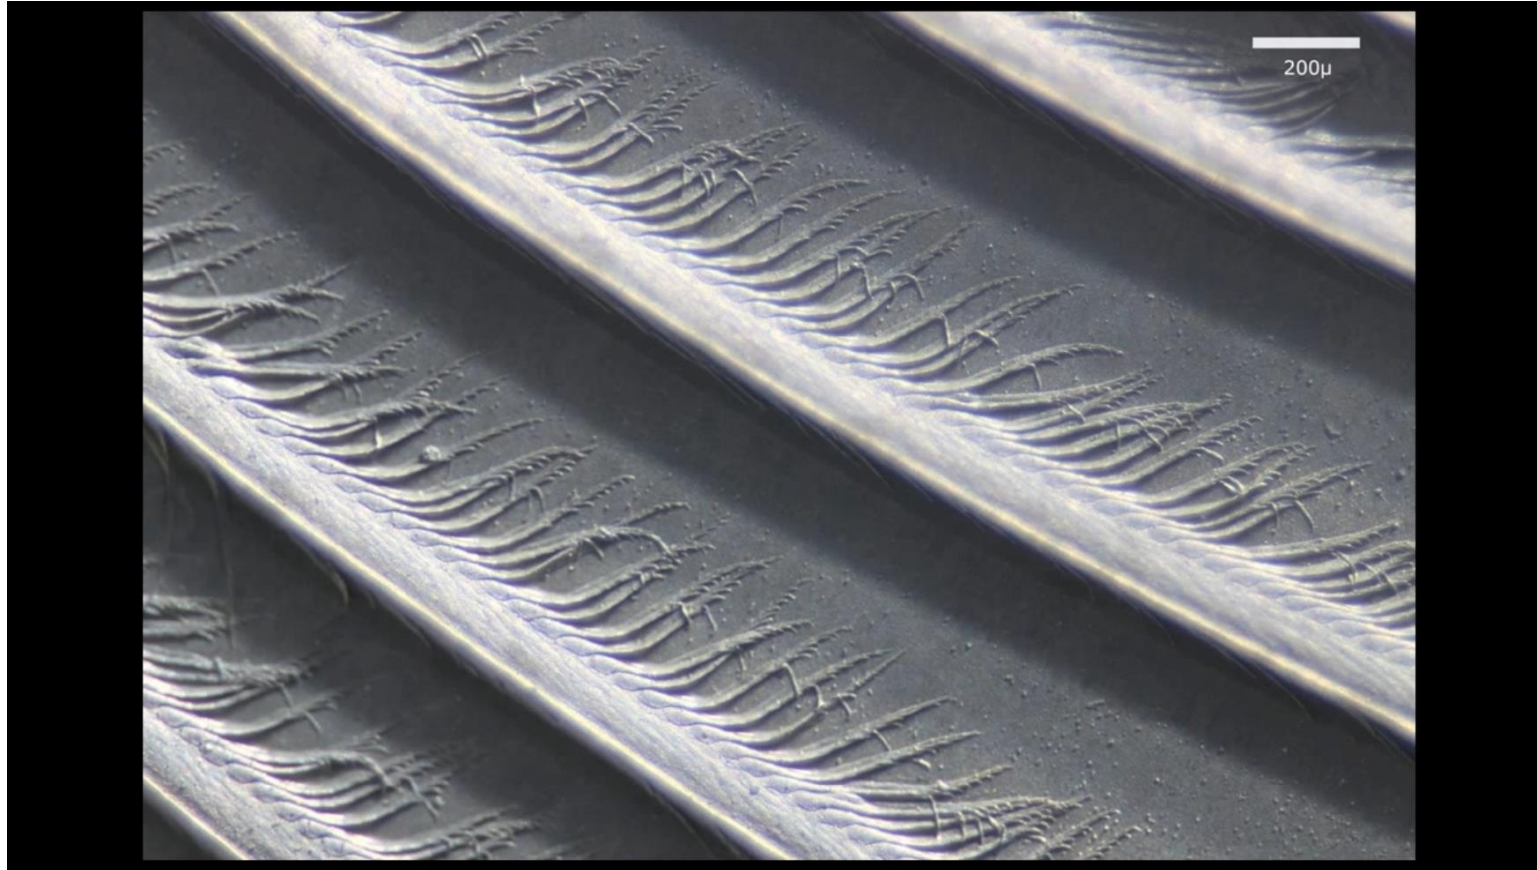


Figure 9: Example geometry measured with the bench and portable configurations. Outer image: rendering under direct lighting. Inset: macro photograph of original sample. Scale shown in upper left. Color images are shown for context and are to similar, but not exact scale.



Sensing Surfaces with GelSight



kimoatmit



138,850 views

<https://www.youtube.com/watch?v=S7gXih4XS7A>

InverseRenderNet: Learning single image inverse rendering

Ye Yu and William A. P. Smith

Department of Computer Science, University of York, UK

{yy1571, william.smith}@york.ac.uk

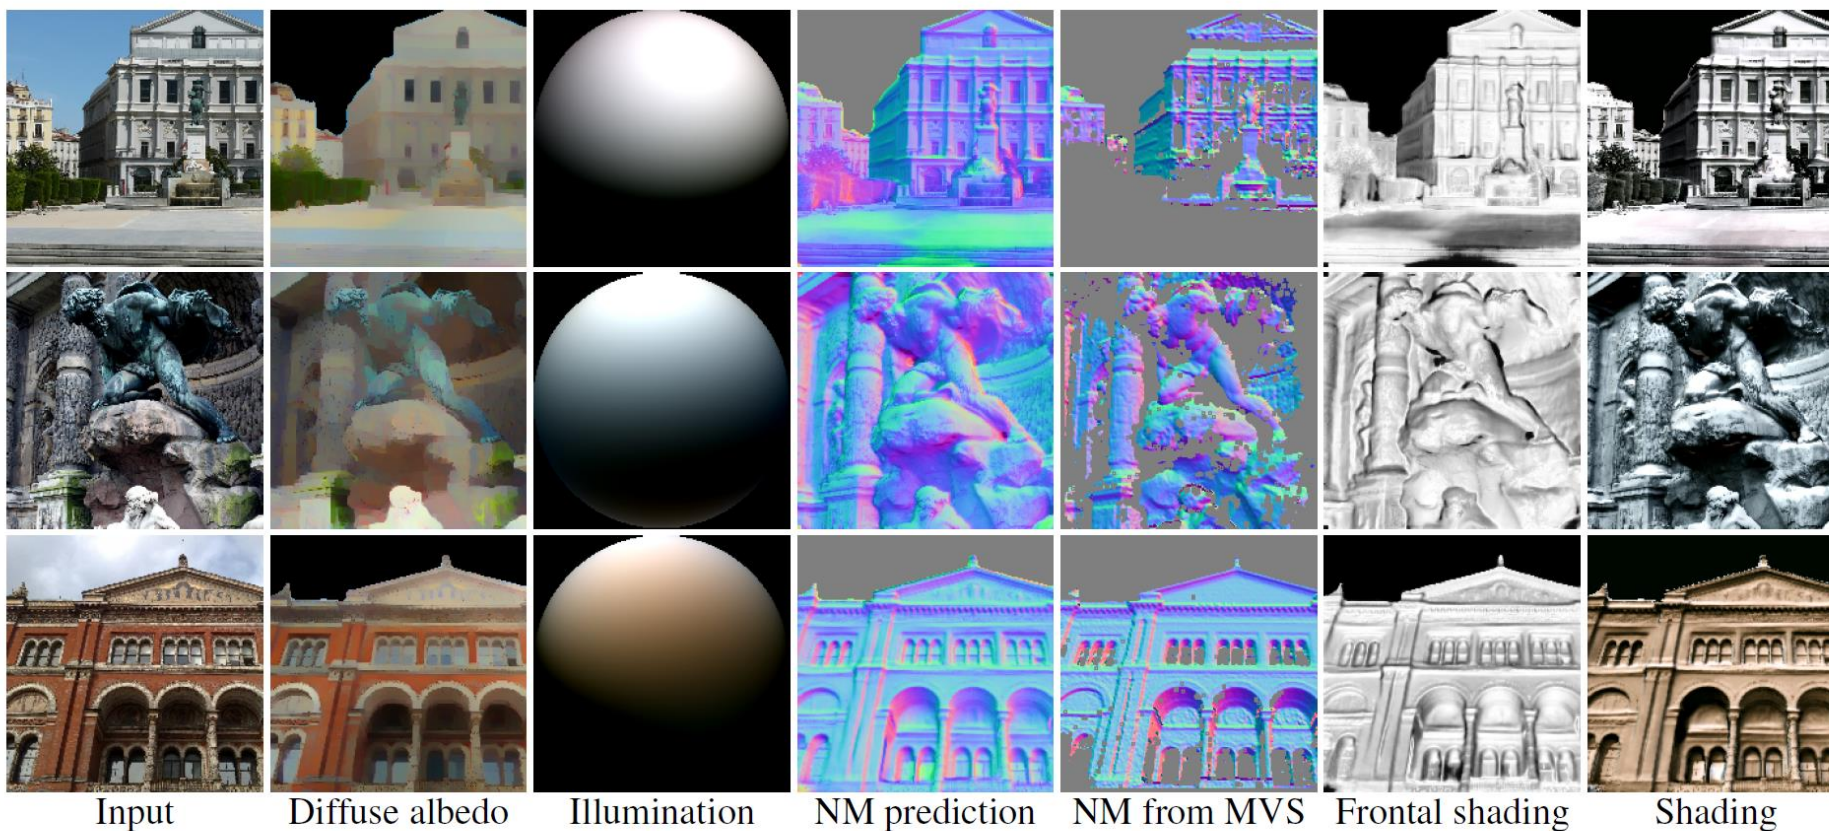


Figure 1: From a single image (col. 1), we estimate albedo and normal maps and illumination (col. 2-4); comparison multi-view stereo result from several hundred images (col. 5); re-rendering of our shape with frontal/estimated lighting (col. 6-7).

Questions?