Visibility Subspaces: Uncalibrated Photometric Stereo with Shadows

Kalyan Sunkavalli, *Harvard University*

Joint work with Todd Zickler and Hanspeter Pfister

Published in the Proceedings of ECCV 2010

http://gvi.seas.harvard.edu/





Shading contains strong perceptual cues about shape

Photometric Stereo

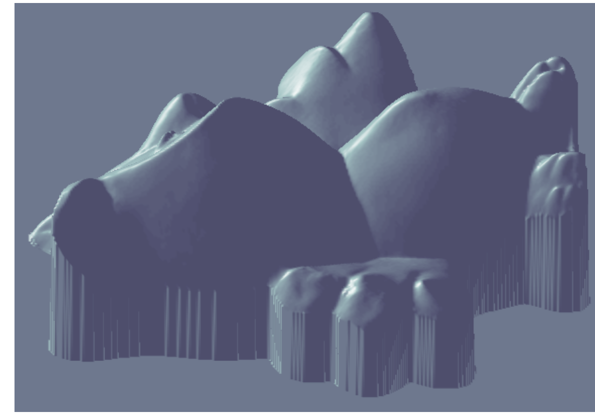

- Use multiple images captured under changing illumination and recover per-pixel surface normals.
- Originally proposed for Lambertian surfaces under directional lighting. Extended to different BRDFs, environment map illumination, etc.
- **One (unavoidable) issue: *how to deal with shadows?***

Lambertian Photometric Stereo

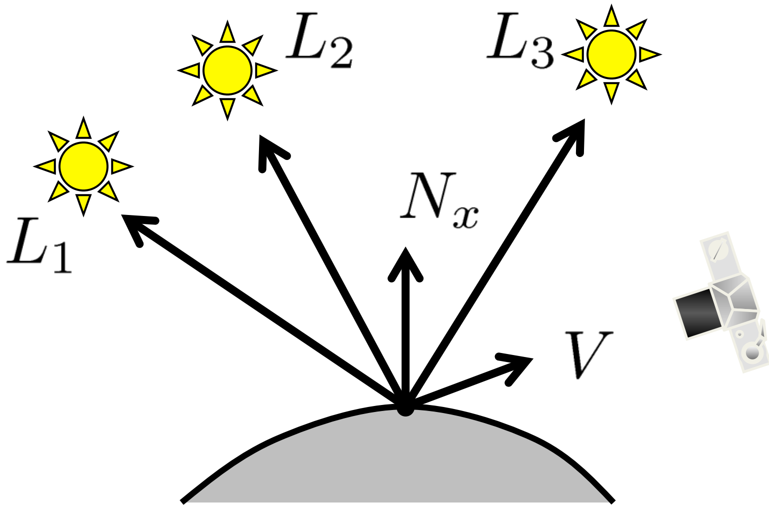

$$I_{x,1} = \rho_x N_x \cdot L_1$$
$$I_{x,2} = \rho_x N_x \cdot L_2$$
$$I_{x,3} = \rho_x N_x \cdot L_3$$

Lambertian Photometric Stereo

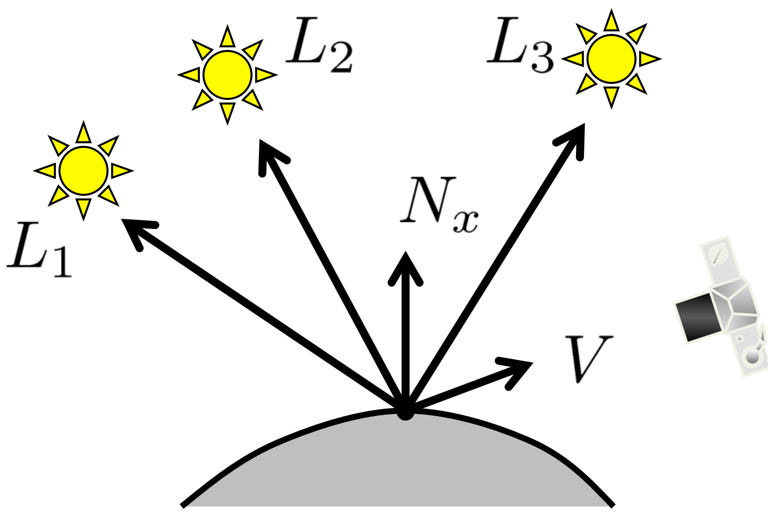

$$I_{x,1} = \rho_x N_x \cdot L_1$$
$$I_{x,2} = \rho_x N_x \cdot L_2$$
$$I_{x,3} = \rho_x N_x \cdot L_3$$

pixels →

lights ↓

$$\begin{bmatrix} I_{1,1} & I_{1,2} & \cdots & I_{1,n} \\ I_{2,1} & I_{2,2} & \cdots & I_{2,n} \\ & & & \\ I_{m,1} & I_{m,2} & \cdots & I_{m,n} \end{bmatrix} = \begin{bmatrix} N_1^T \\ N_2^T \\ \\ N_m^T \end{bmatrix} \begin{bmatrix} L_1 & L_2 & \cdots & L_n \end{bmatrix}$$

$$\mathbf{I}_{m\times n} = \mathbf{N}^T_{m\times 3}\,\mathbf{L}_{3\times n}$$

Lambertian Photometric Stereo

$$\mathbf{I}_{m \times n} \quad = \quad \mathbf{N}^T_{m \times 3} \, \mathbf{L}_{3 \times n}$$

- Images of a Lambertian surface under directional lighting form a Rank-3 matrix. [Shashua '97]

Lambertian Photometric Stereo

$$\mathbf{I}_{m\times n} \quad = \quad \mathbf{N}^T_{m\times 3}\,\mathbf{L}_{3\times n}$$

- Images of a Lambertian surface under directional lighting form a Rank-3 matrix.
- Photometric Stereo (*calibrated* lighting)

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

[Woodham '78, Silver '80]

Lambertian Photometric Stereo

$$\mathbf{I}_{m\times n} \quad = \quad \mathbf{N}^T_{m\times 3}\,\mathbf{L}_{3\times n}$$

- Images of a Lambertian surface under directional lighting form a Rank-3 matrix.

- Photometric Stereo (*uncalibrated* lighting)

$$\mathbf{I} = \mathbf{U}\boldsymbol{\Sigma}_{3\times 3}\mathbf{V}^T$$

$$\mathbf{N^T} = \mathbf{U}(\boldsymbol{\Sigma}_{3\times 3})^{1/2}\;,\;\mathbf{L} = (\boldsymbol{\Sigma}_{3\times 3})^{1/2}\mathbf{V}^T$$

Ambiguity $\mathbf{I} = \mathbf{N}^T\mathbf{L} = \mathbf{N}^T\mathbf{A}^{-1}\mathbf{A}\mathbf{L}$

[Hayakawa ’94, Epstein et al. ’96, Belhumeur et al. ‘99]

Shadows in Photometric Stereo

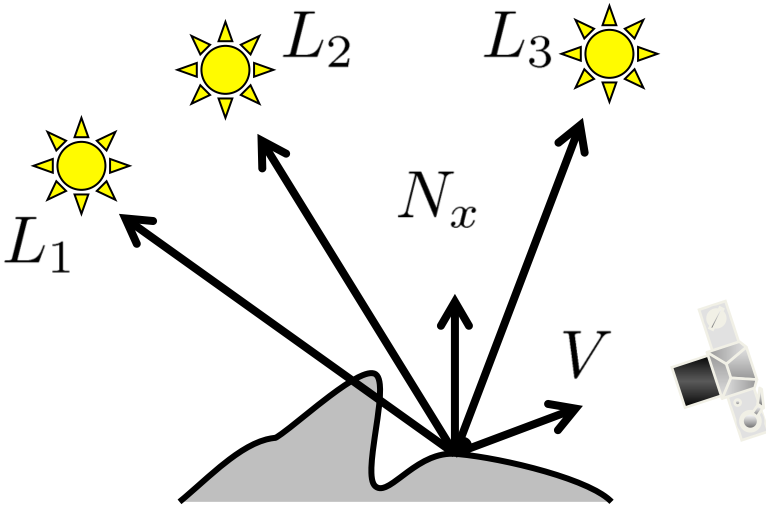

$$I_{x,1} = \rho_x N_x \cdot (V_{x,1} L_1)$$
$$I_{x,2} = \rho_x N_x \cdot (V_{x,2} L_2)$$
$$I_{x,3} = \rho_x N_x \cdot (V_{x,3} L_3)$$

$$\mathbf{I}_{m\times n} = \left(\mathbf{N}^T_{m\times 3}\, \mathbf{L}_{3\times n}\right) \otimes \mathbf{V}_{m\times n}$$

Shadows in Photometric Stereo

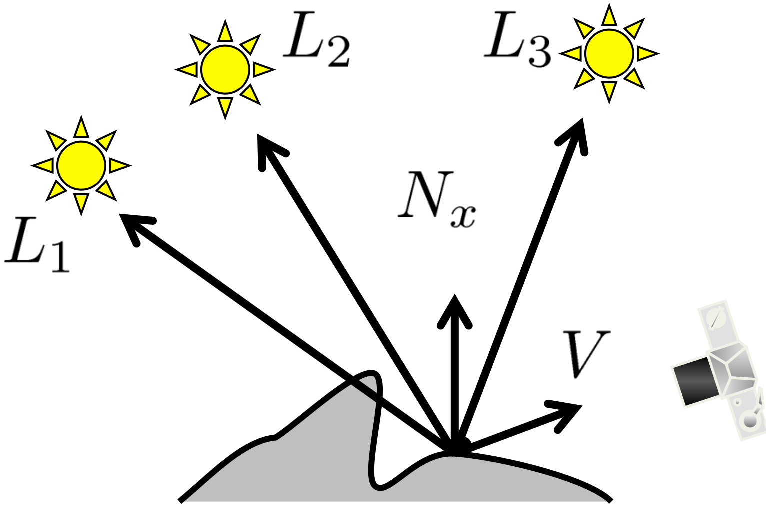

$$I_{x,1} = \rho_x N_x \cdot (V_{x,1} L_1)$$
$$I_{x,2} = \rho_x N_x \cdot (V_{x,2} L_2)$$
$$I_{x,3} = \rho_x N_x \cdot (V_{x,3} L_3)$$

$$\mathbf{I}_{m\times n} = \ (\mathbf{N}^T_{m\times 3}\ \mathbf{L}_{3\times n}) \otimes \mathbf{V}_{m\times n}$$

- Photometric Stereo (*calibrated* lighting)

  $$N_i^T = ((I_i \otimes V_i)(\mathbf{L} \otimes V_i)^T)((\mathbf{L} \otimes V_i)(\mathbf{L} \otimes V_i)^T)^{-1}$$

- Photometric Stereo (*uncalibrated* lighting)

  $\mathbf{I} \otimes \mathbf{V} = (\mathbf{N}^T\ \mathbf{L}) \otimes \mathbf{V}$ (factorization with missing data)

Shadows in Photometric Stereo

- Previous work: Detect shadowed pixels and discard them.

- Intensity-based thresholding
  - Threshold requires (unknown) albedo
- Use calibrated lights to estimate shadows
  [Coleman & Jain ‘82, Chandraker & Kriegman ’07]
- Smoothness constraints on shadows
  [Chandraker & Kriegman ’07, Hernandez et al. ’08]
- Use many (100s) images.
  [Wu et al. ’06, Wu et al. 10]

Shadows in Photometric Stereo

- Our work analyzes the effect of shadows on scene appearance.

- We show that shadowing leads to distinct appearance subspaces.

- This results in:
  - A novel bound on the dimensionality of (Lambertian) scene appearance.
  - An uncalibrated Photometric Stereo algorithm that works in the presence of shadows.

Shadows and Scene Appearance

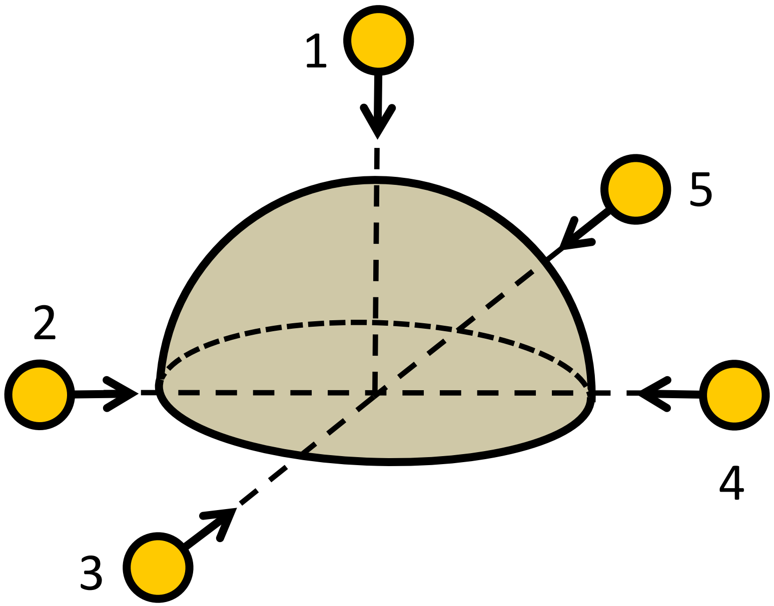

Scene

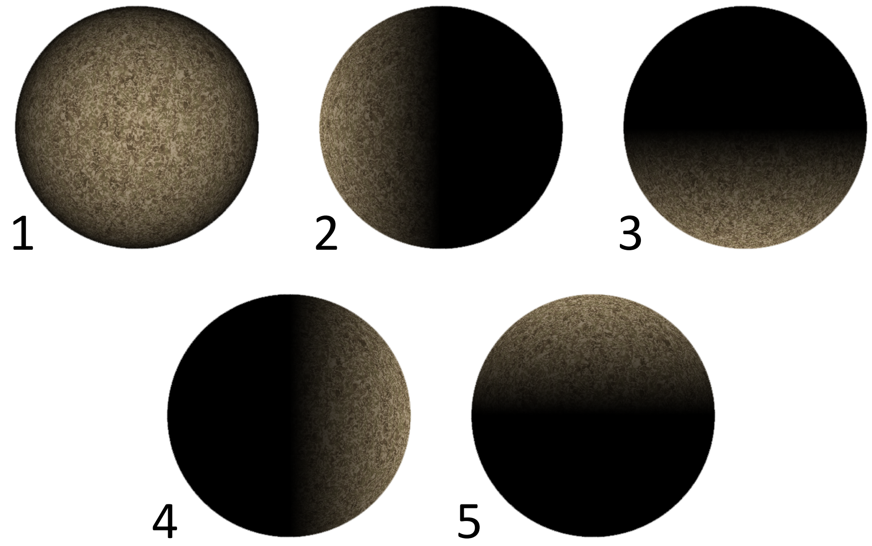

Images

Shadows and Scene Appearance

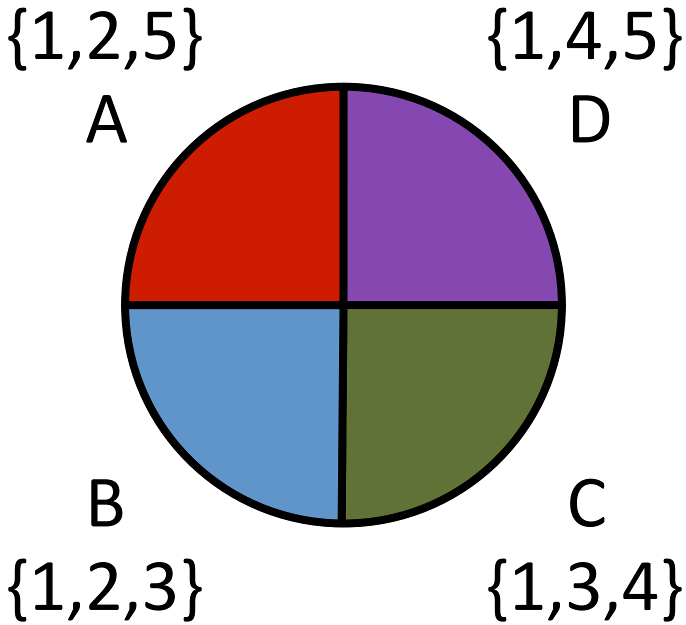

Visibility Regions

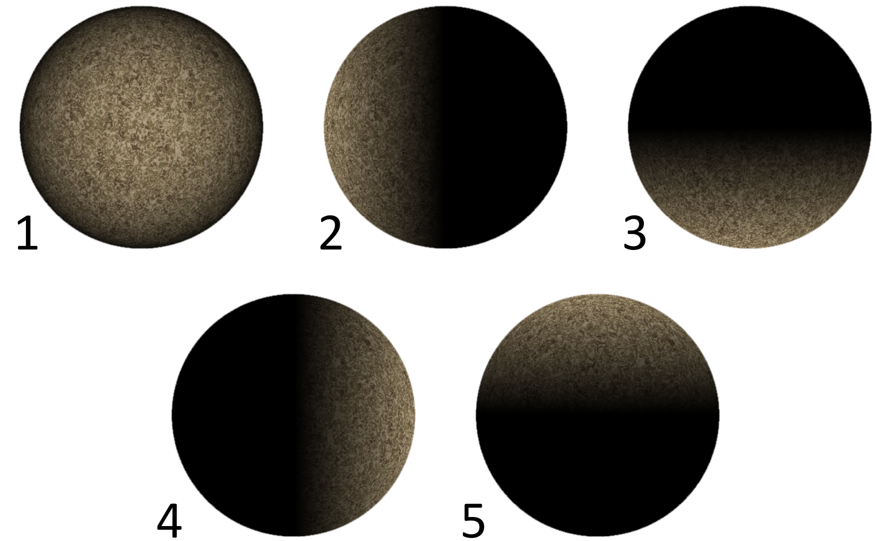

Images

Shadows and Scene Appearance

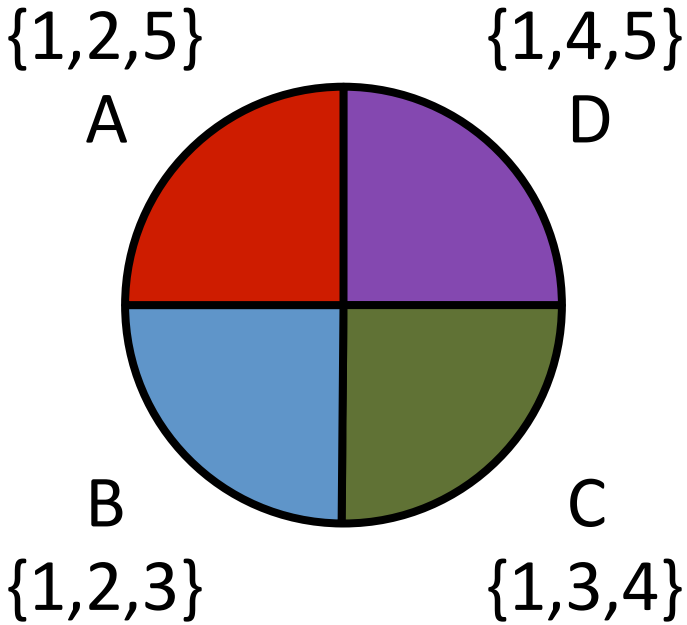

Visibility Regions

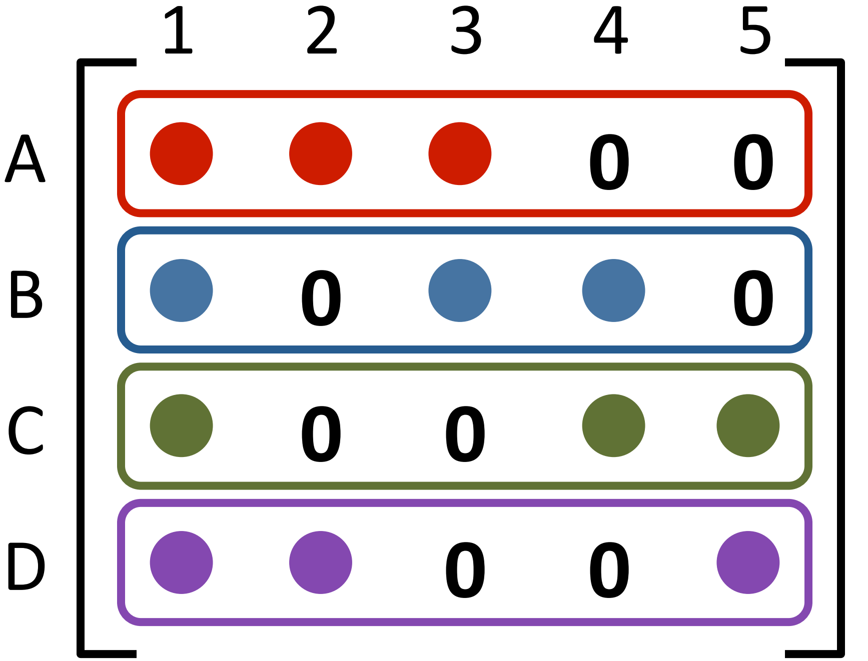

Image Matrix
Rank-5

Shadows and Scene Appearance

Lambertian points lit
by directional lights
Rank-3 submatrix

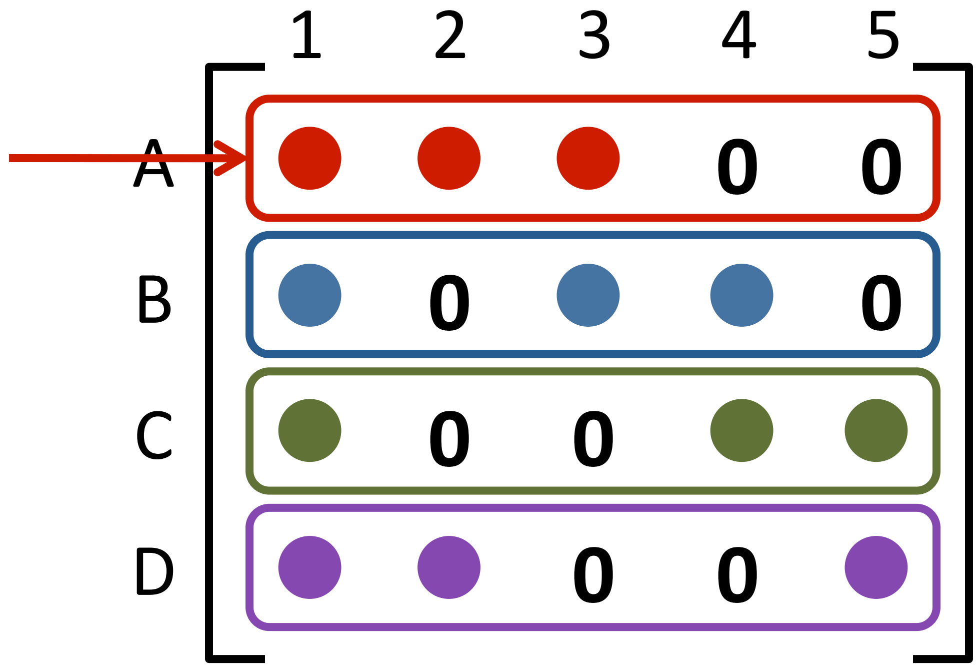

Image Matrix
Rank-5

Shadows and Scene Appearance

Different Rank-3 submatrices

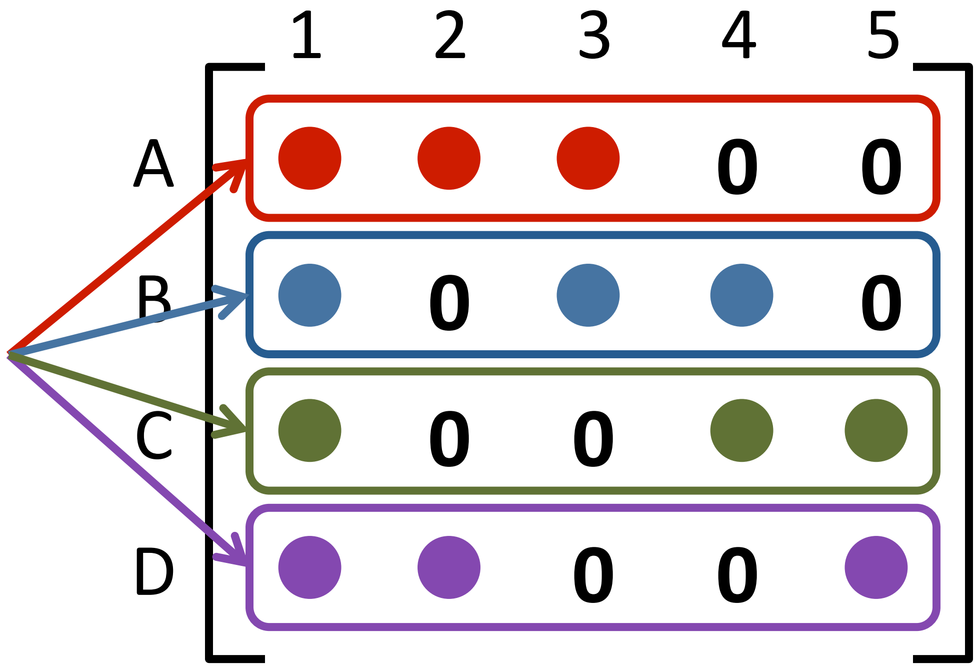

Image Matrix
Rank-5

Visibility Subspaces

Scene points with same visibility

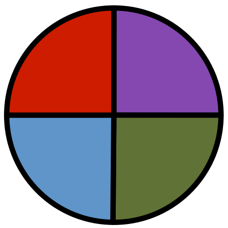

Rank-3 subspaces of image matrix

Visibility Subspaces

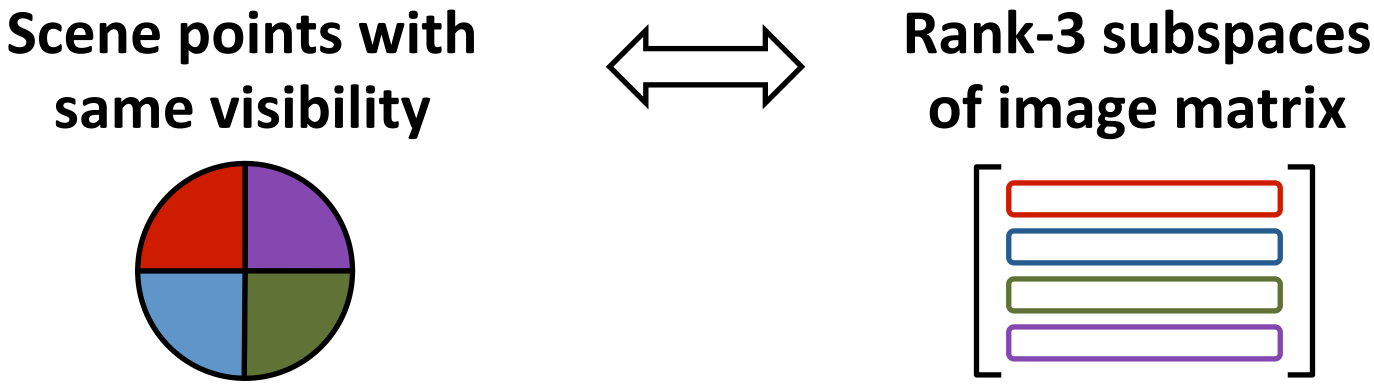

1. *Dimensionality of scene appearance with (cast) shadows:*
   Images of a *Lambertian* scene illuminated by any combination of ***n*** light sources lie in a linear space with dimension at most $\boldsymbol{3(2^n)}$.

   Previous work excludes analysis of cast shadows.

Visibility Subspaces

Scene points with same visibility ⟺ **Rank-3 subspaces of image matrix**

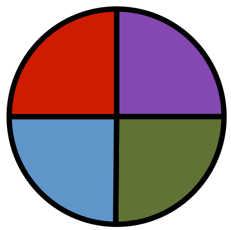

1. *Dimensionality of scene appearance with (cast) shadows*

2. Visibility regions can be recovered through subspace estimation (leading to an uncalibrated Photometric Stereo algorithm).

Estimating Visibility Subspaces

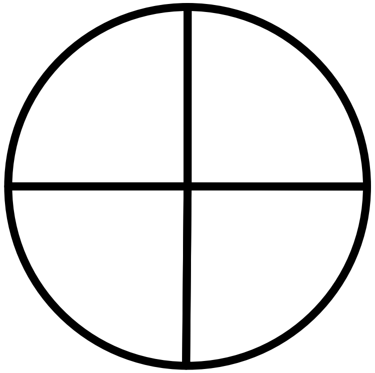

- Find visibility regions by looking for Rank-3 subspaces (using RANSAC-based subspace estimation).

Estimating Visibility Subspaces

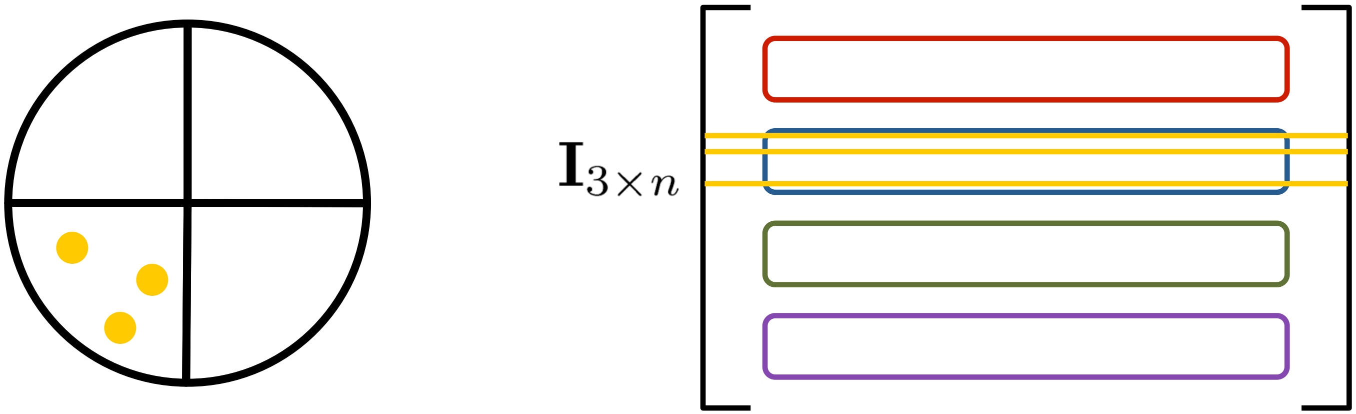

- Sample 3 points and construct lighting basis $\mathbf{L}$ from the image intensities:

$$\mathbf{I}_{3\times n} = \mathbf{N}^T_{3\times 3}\mathbf{L}_{3\times n}$$

Estimating Visibility Subspaces

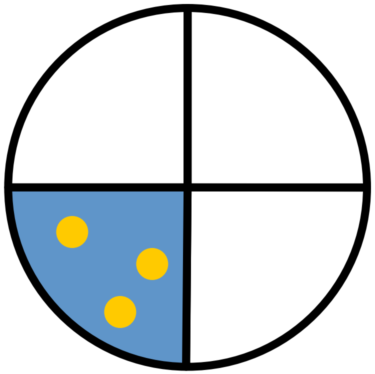

$\mathbf{I}_{3\times n}$


- Sample 3 points and construct lighting basis $\mathbf{L}$ from the image intensities:

$$\mathbf{I}_{3\times n} = \mathbf{N}^T_{3\times 3}\mathbf{L}_{3\times n}$$

- If points are in *same* visibility subspace, $\mathbf{L}$ is a valid basis for *entire subspace*.

Estimating Visibility Subspaces

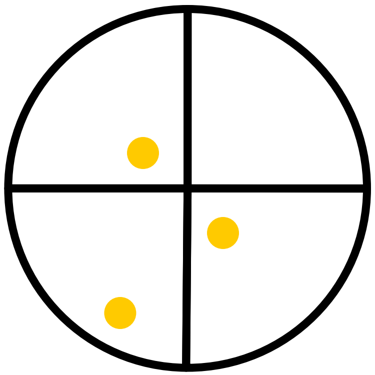

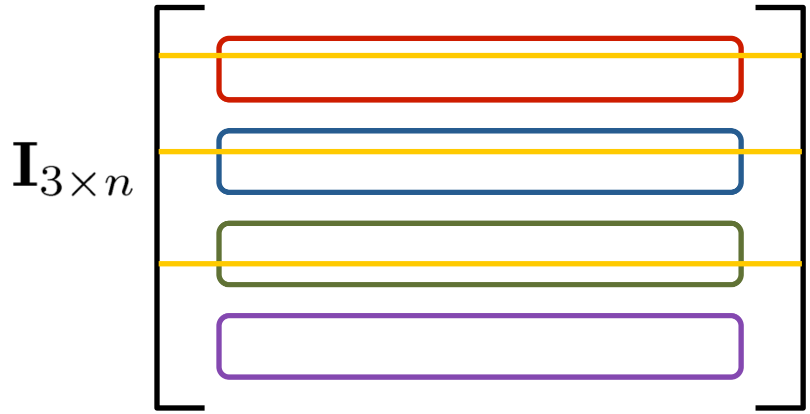

- Sample 3 points and construct lighting basis $\mathbf{L}$ from the image intensities:

$$\mathbf{I}_{3\times n} = \mathbf{N}^T_{3\times 3}\mathbf{L}_{3\times n}$$

- If points are in *same* visibility subspace, $\mathbf{L}$ is a valid basis for *entire subspace*.

- If not, $\mathbf{L}$ is not a valid basis for *any subspace*.

Estimating Visibility Subspaces

1. Sample 3 points in scene and construct lighting basis from their image intensities: $\mathbf{I}_3 = \hat{\mathbf{N}}_3^T \hat{\mathbf{L}}_3$

2. Compute normals at all points using this basis: $\hat{N}_i^T = I_i(\hat{\mathbf{L}}_3)^+$.

3. Compute error of this basis: $E_i = ||I_i - \hat{N}_i^T \hat{\mathbf{L}}_3||^2$.

4. Mark points with error $E_i < \epsilon$ as inliers.

5. Repeat 1-4 and mark largest inlier-set found as subspace with lighting basis $\hat{\mathbf{L}}_k$. Remove inliers from pixel-set.

6. Repeat 1-5 until all visibility subspaces have been recovered.

Estimating Visibility Subspaces

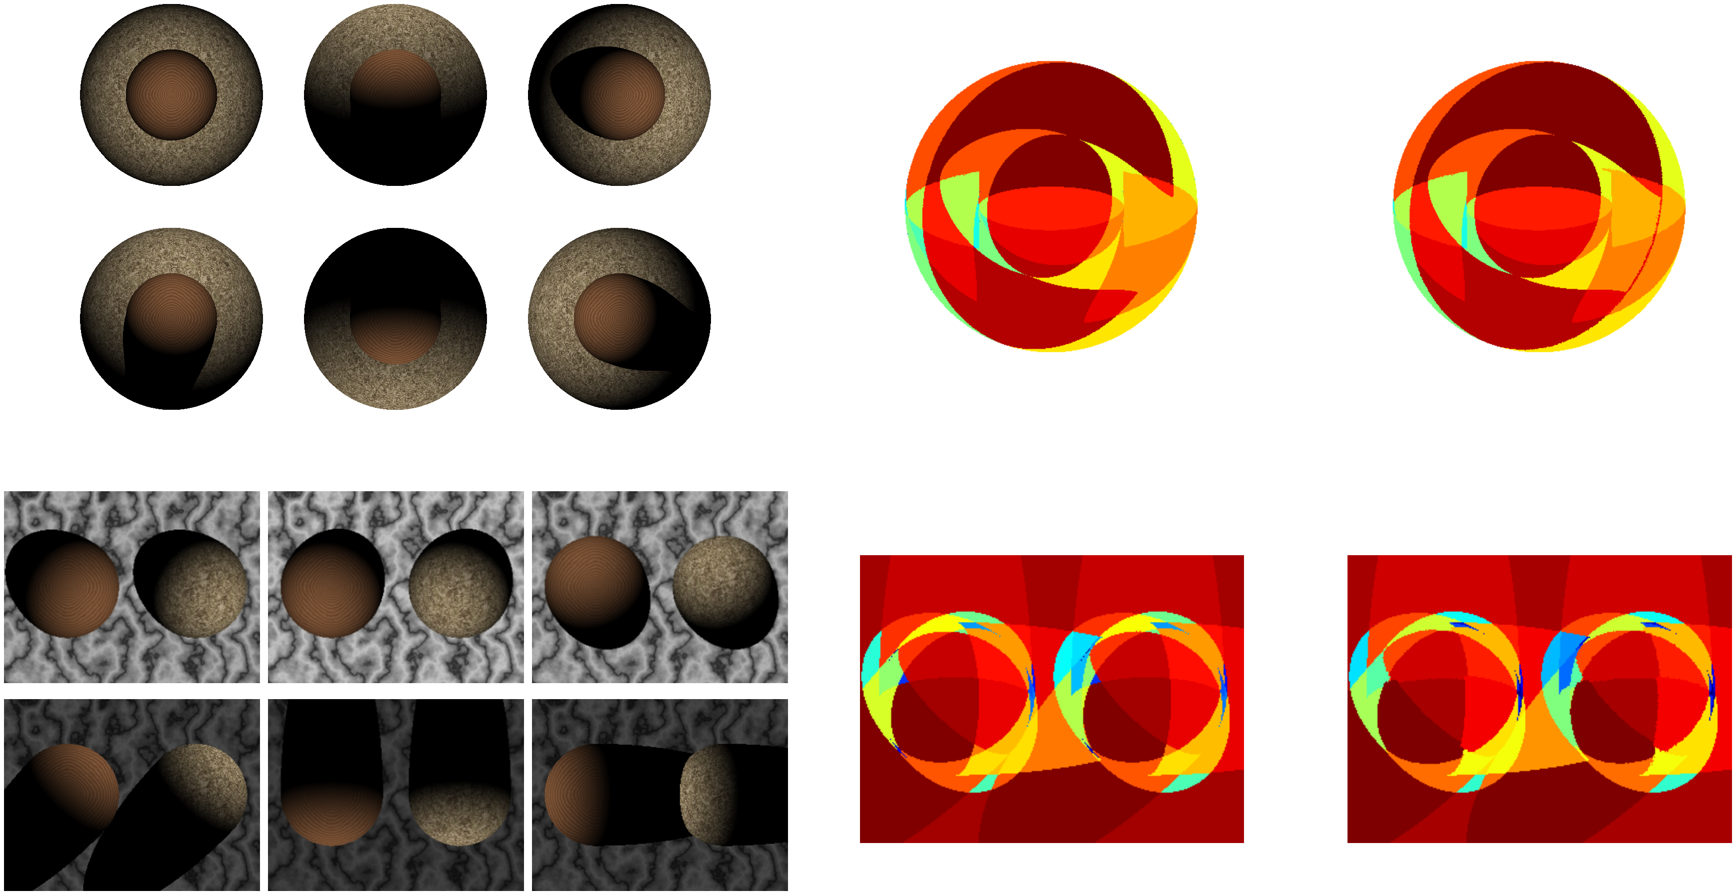

Images

True subspaces

Estimated subspaces

Subspaces to Surface normals

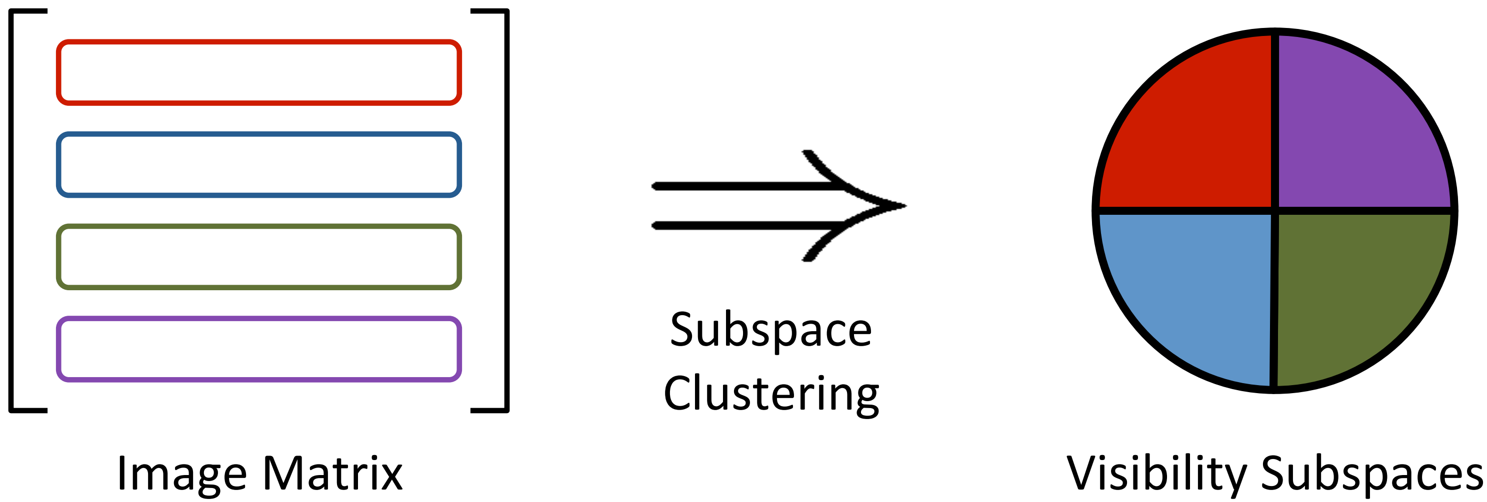

- Subspace clustering gives us a labeling of the scene points into regions with same visibility.

- Can we figure out the true visibility (and surface normals) from this?

Subspaces to Surface normals

- Subspace clustering recovers normals and lights:

$$\mathbf{I}_k = \underbrace{\hat{\mathbf{N}}_k^T}_{\text{Subspace Normals}} \underbrace{\hat{\mathbf{L}}_k}_{\text{Subspace Lights}}$$

Subspaces to Surface normals

- Subspace clustering recovers normals and lights:

$$\mathbf{I}_k = \hat{\mathbf{N}}_k^T \hat{\mathbf{L}}_k$$

- There is a 3X3 linear ambiguity in these normals and lights:

$$\hat{\mathbf{N}}_k^T = \underbrace{\mathbf{N}_k^T}_{\text{True Normals}} \underbrace{\mathbf{A}_k^{-1}}_{\text{Subspace Ambiguity}} \;,\; \hat{\mathbf{L}}_k = \underbrace{\mathbf{A}_k}_{\text{Subspace Ambiguity}} \underbrace{\mathbf{L}_k}_{\text{True Lights}}$$

True Normals

Subspace Ambiguity

True Lights

Subspaces to Surface normals

- Subspace clustering recovers normals and lights:

$$\mathbf{I}_k = \hat{\mathbf{N}}_k^T \hat{\mathbf{L}}_k$$

- There is a 3X3 linear ambiguity in these normals and lights:

$$\hat{\mathbf{N}}_k^T = \underbrace{\mathbf{N}_k^T}_{\text{True Normals}} \underbrace{\mathbf{A}_k^{-1}}_{\text{Subspace Ambiguity}}, \quad \hat{\mathbf{L}}_k = \underbrace{\mathbf{A}_k}_{\text{Subspace Ambiguity}} \underbrace{\mathbf{L}_k}_{\text{True Lights}}$$

True Normals | Subspace Ambiguity | True Lights

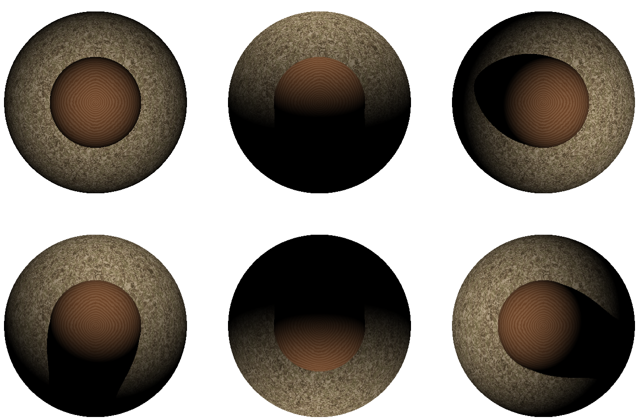

Images

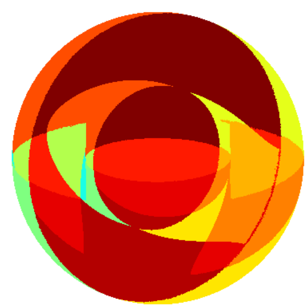

Estimated subspaces

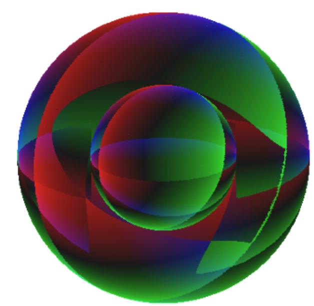

Subspace normals

Subspaces to Surface normals

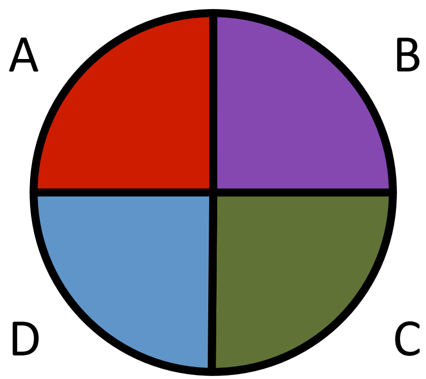

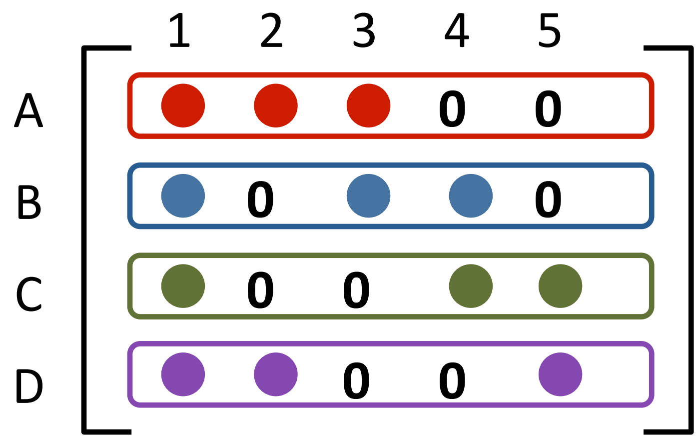

$$
\begin{aligned}
\mathbf{L}_A &= \begin{bmatrix} L_1 & L_2 & L_3 & 0 & 0 \end{bmatrix} \\
\mathbf{L}_B &= \begin{bmatrix} L_1 & 0 & L_3 & L_4 & 0 \end{bmatrix} \\
\mathbf{L}_C &= \begin{bmatrix} L_1 & L_0 & 0 & L_4 & L_5 \end{bmatrix} \\
\mathbf{L}_D &= \begin{bmatrix} L_1 & L_2 & 0 & 0 & L_5 \end{bmatrix}
\end{aligned}
$$

Subspaces to Surface normals

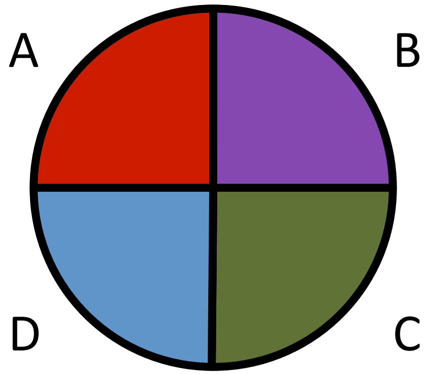

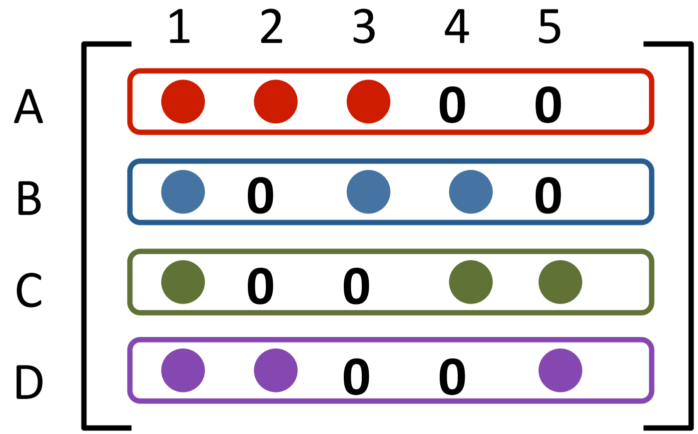

$$\mathbf{L}_k \quad = \quad \mathbf{L} \otimes V_k$$

Subspaces to Surface normals

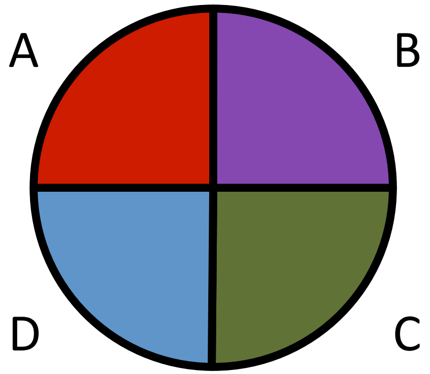

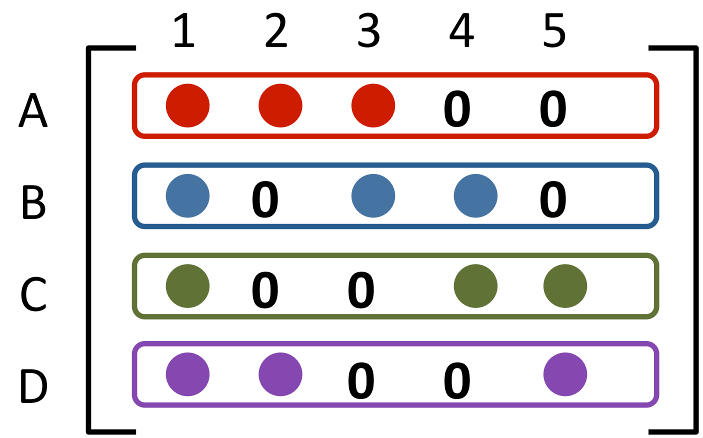

$$\mathbf{L}_k = \mathbf{L} \otimes V_k$$

$$\mathbf{A}_k^{-1}\hat{\mathbf{L}}_k = \mathbf{L} \otimes V_k$$

$\mathbf{A}_k^{-1}$: Subspace ambiguity

$\hat{\mathbf{L}}_k$: Subspace light basis (from clustering)

$\mathbf{L}$: True lights

V_k: Visibility

Subspaces to Surface normals

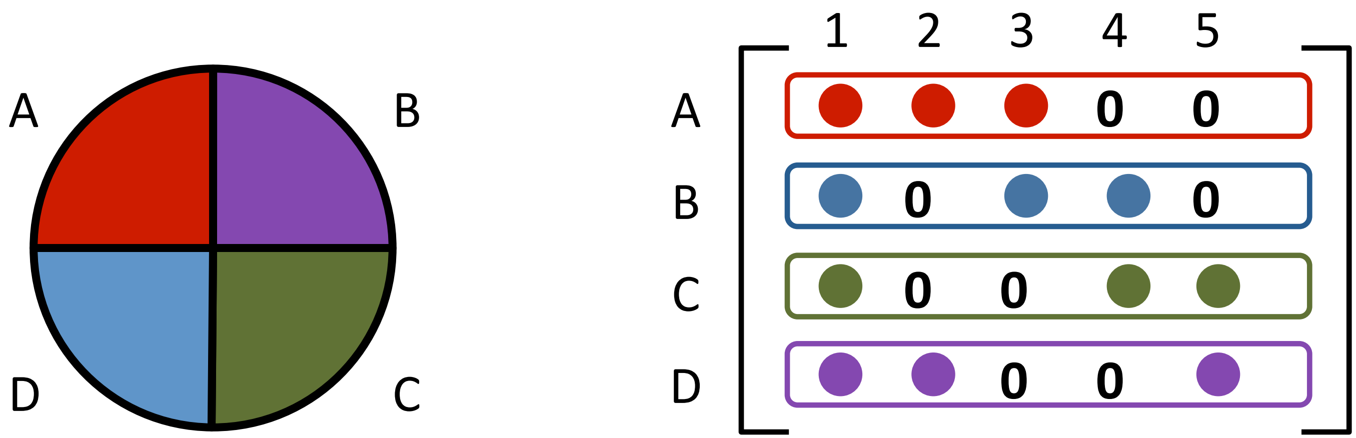

$$\mathbf{I}_A = \mathbf{N}_A^T \left[\begin{array}{ccccc} L_1 & L_2 & L_3 & 0 & 0 \end{array} \right]$$

$$\approx \hat{\mathbf{N}}_A^T \left[\begin{array}{ccccc} \hat{L}_1 & \hat{L}_2 & \hat{L}_3 & \hat{L}_4 & \hat{L}_5 \end{array} \right]$$

$$\mathbf{I}_{A,4} \approx 0 \quad \Rightarrow \quad ||\hat{L}_4|| \approx 0$$

$$\mathbf{I}_{A,5} \approx 0 \quad \Rightarrow \quad ||\hat{L}_5|| \approx 0$$

Subspaces to Surface normals

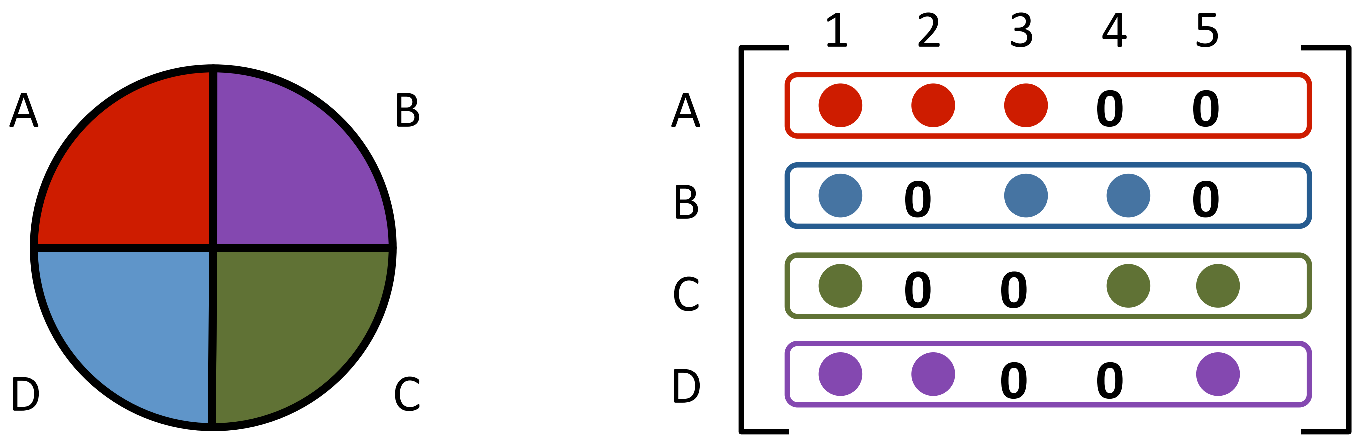

$$\mathbf{I}_A = \mathbf{N}_A^T \left[\; L_1 \quad L_2 \quad L_3 \quad 0 \quad 0 \; \right]$$
$$\approx \hat{\mathbf{N}}_A^T \left[\; \hat{L}_1 \quad \hat{L}_2 \quad \hat{L}_3 \quad \hat{L}_4 \quad \hat{L}_5 \; \right]$$

$$V_{k,j} = ||\hat{\mathbf{L}}_{k,j}|| > \tau$$

Visibility of subspace

Magnitude of subspace light basis

independent of scene properties

Subspaces to Surface normals

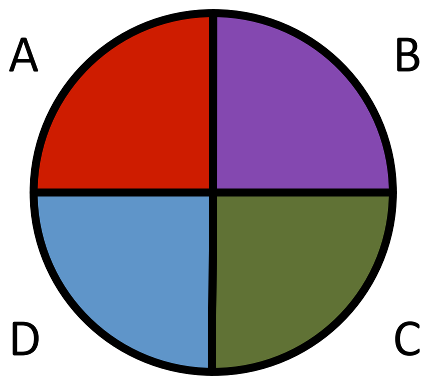

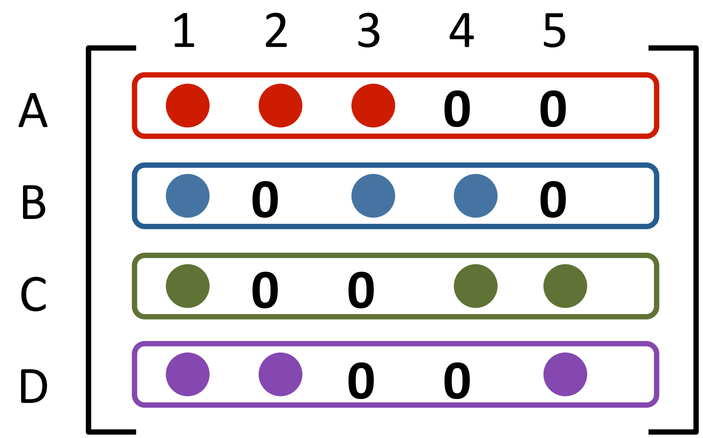

$$\mathbf{A}_k^{-1}\hat{\mathbf{L}}_k = \mathbf{L} \otimes V_k$$

Subspace ambiguity (**unknown**)

Subspace light basis (from clustering)

True lights (**unknown**)

Visibility (computed from subspace lighting)

Subspaces to Surface normals

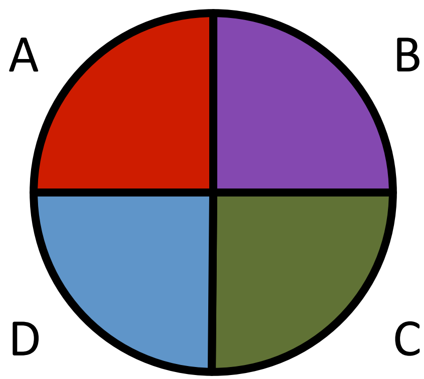

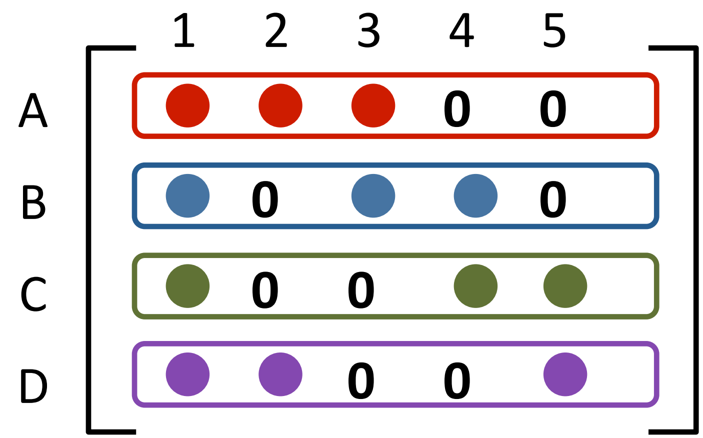

$$\mathbf{A}_k^{-1}\hat{\mathbf{L}}_k = \mathbf{L} \otimes V_k$$

- Linear system of equations
- Solve for ambiguities and true light sources
- Avoid trivial solution ($\mathbf{A}_k^{-1} = 0, V_k = 0$) by setting $\mathbf{A}_s^{-1} = \mathbf{I}$
- Transform subspace normals by estimated ambiguities

Subspaces to Surface normals

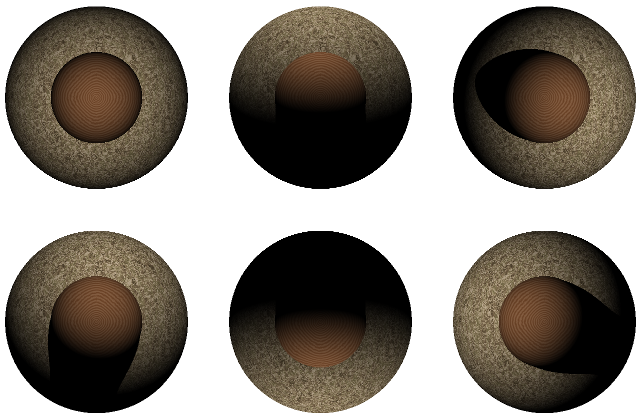

Images

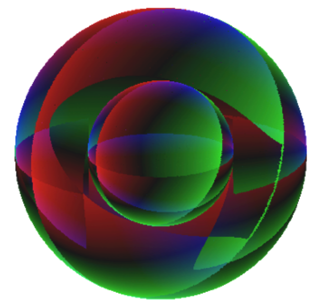

Subspace normals

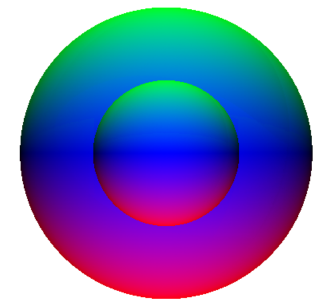

Transformed normals

Visibility Subspaces

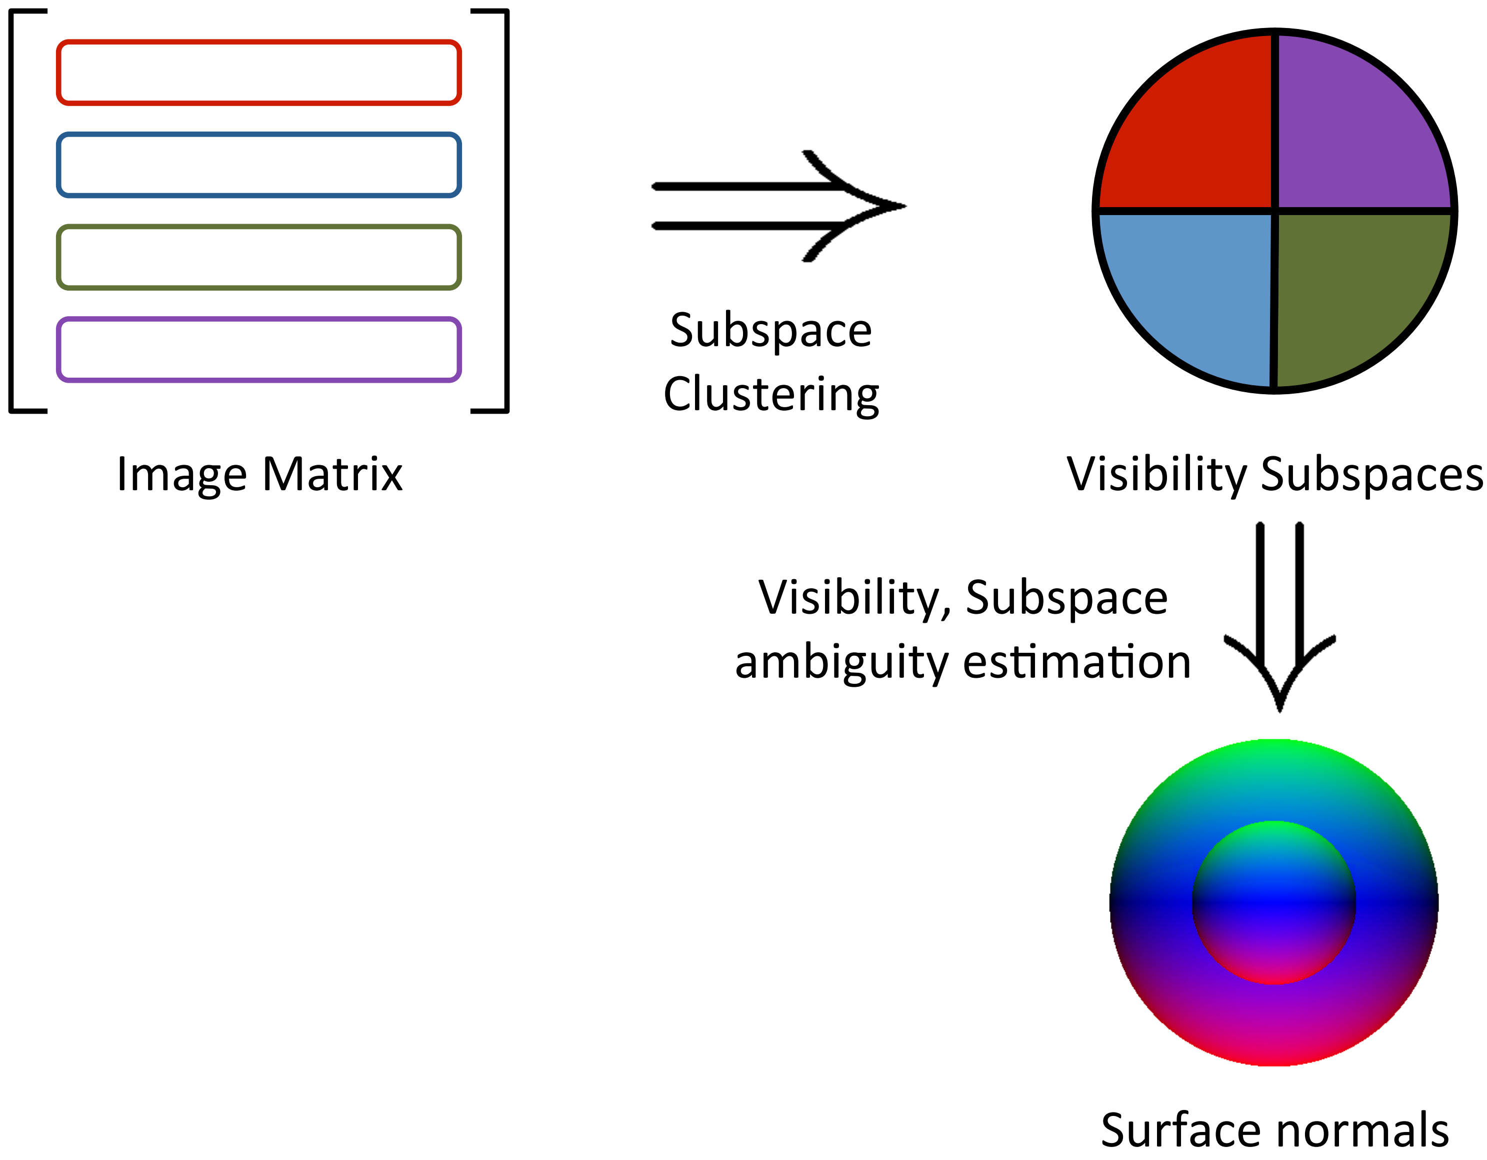

Results (synthetic data)

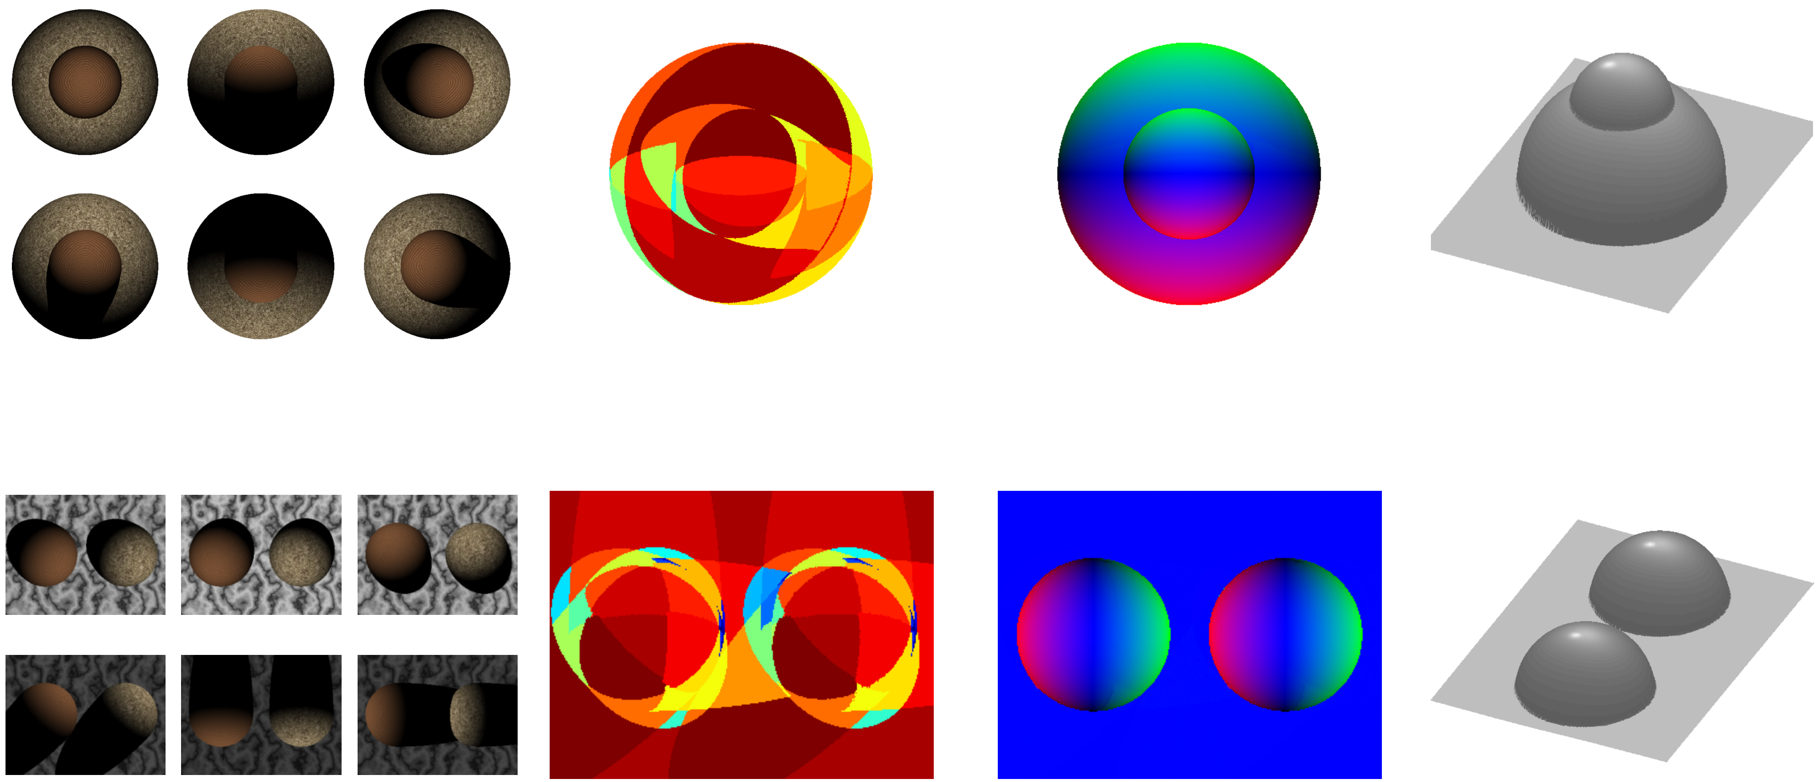

Images | Estimated subspaces | Estimated normals | Estimated depth

Results (captured data)

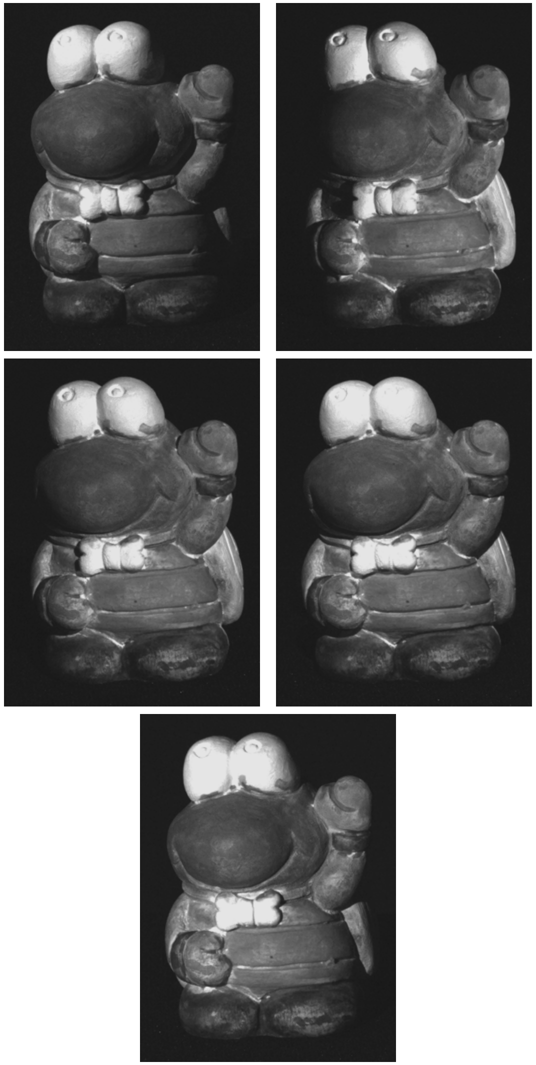

5 Images

Estimated subspaces

Estimated normals

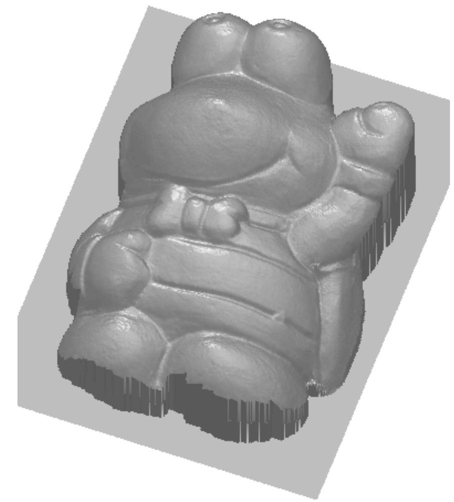

Estimated depth

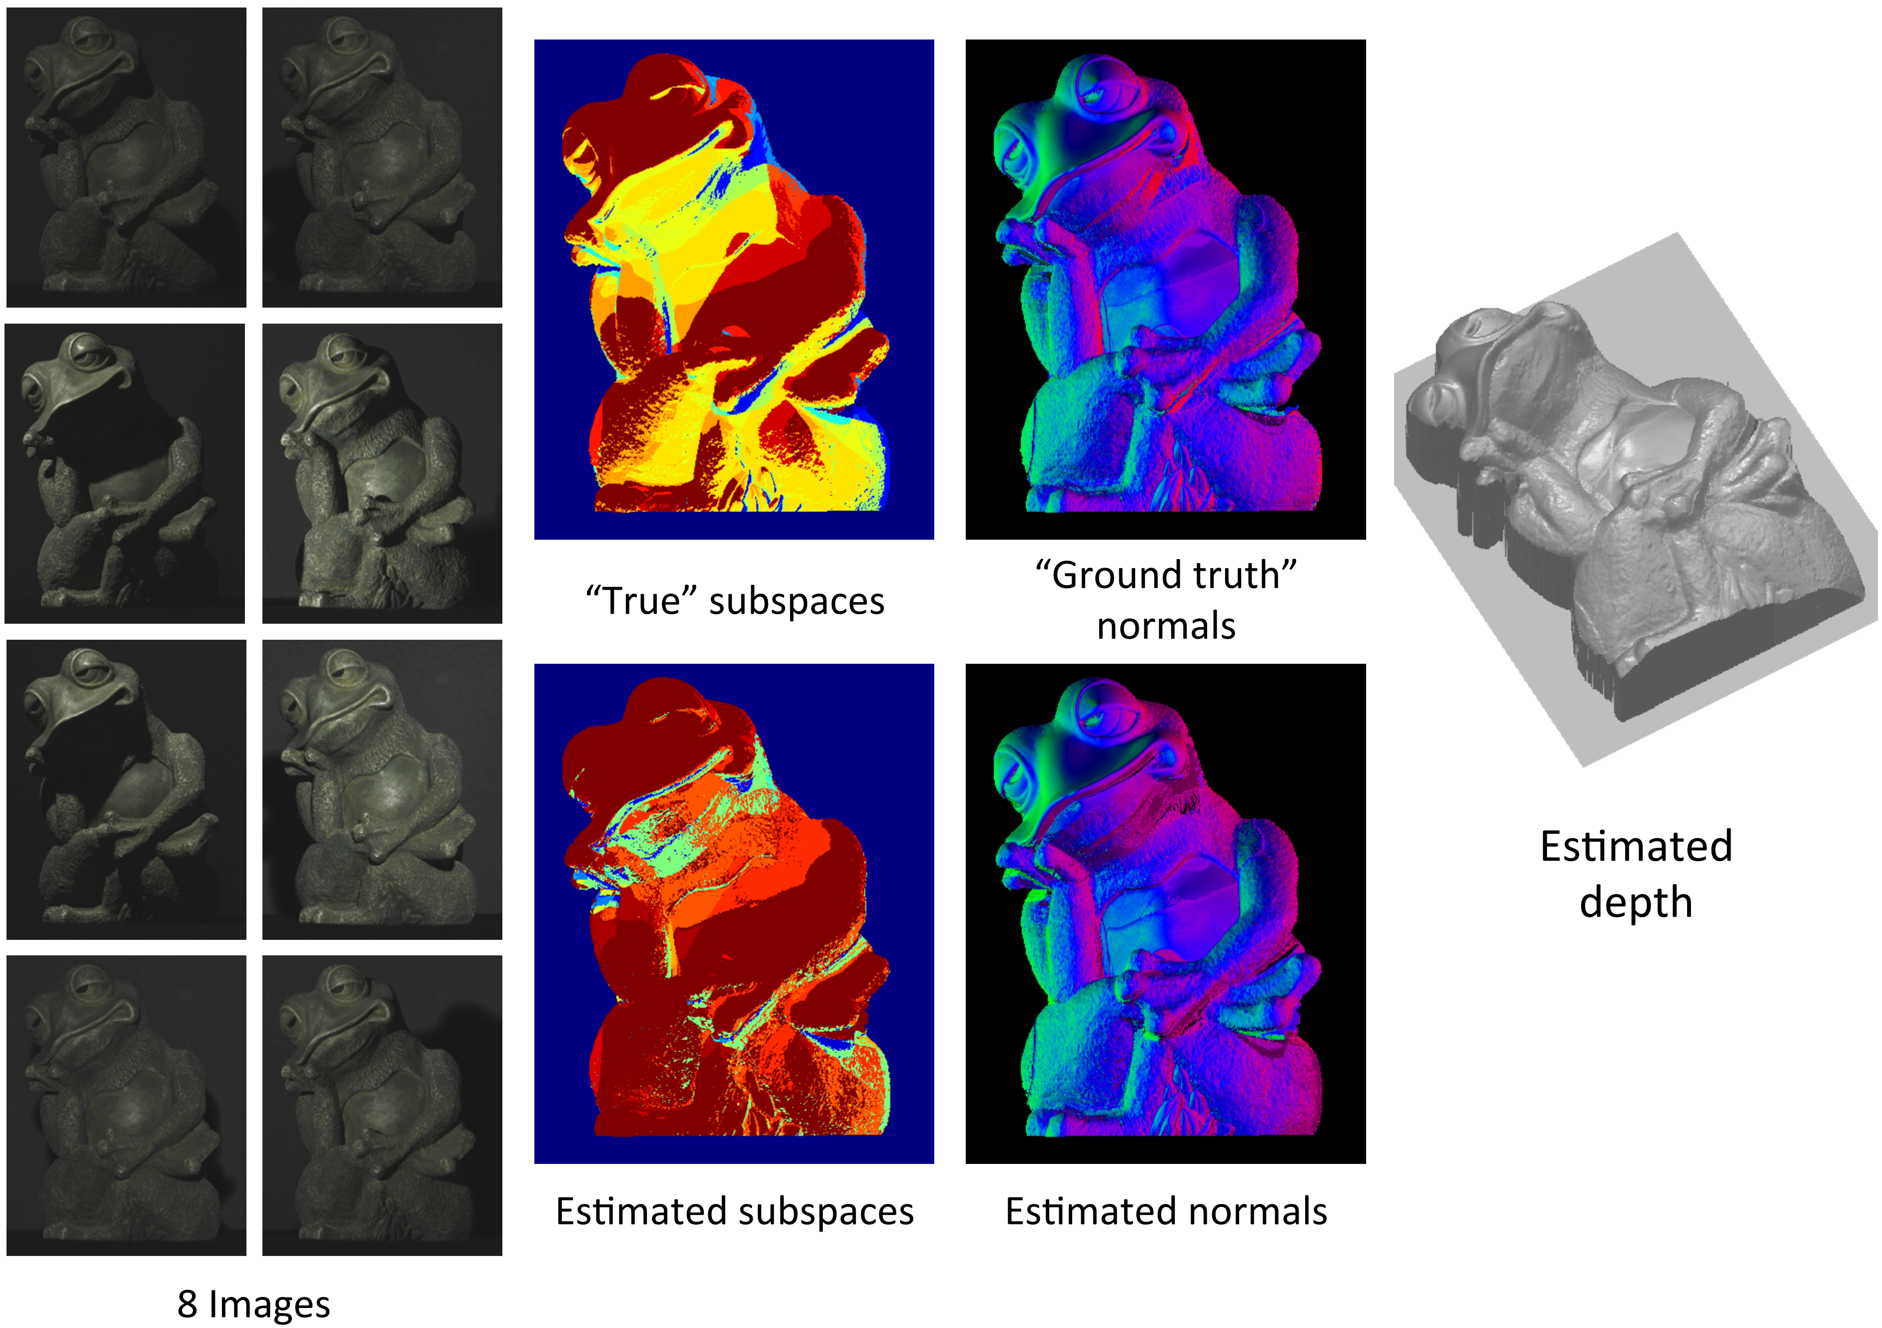
"True" subspaces
"Ground truth" normals
Estimated depth
Estimated subspaces
Estimated normals
8 Images

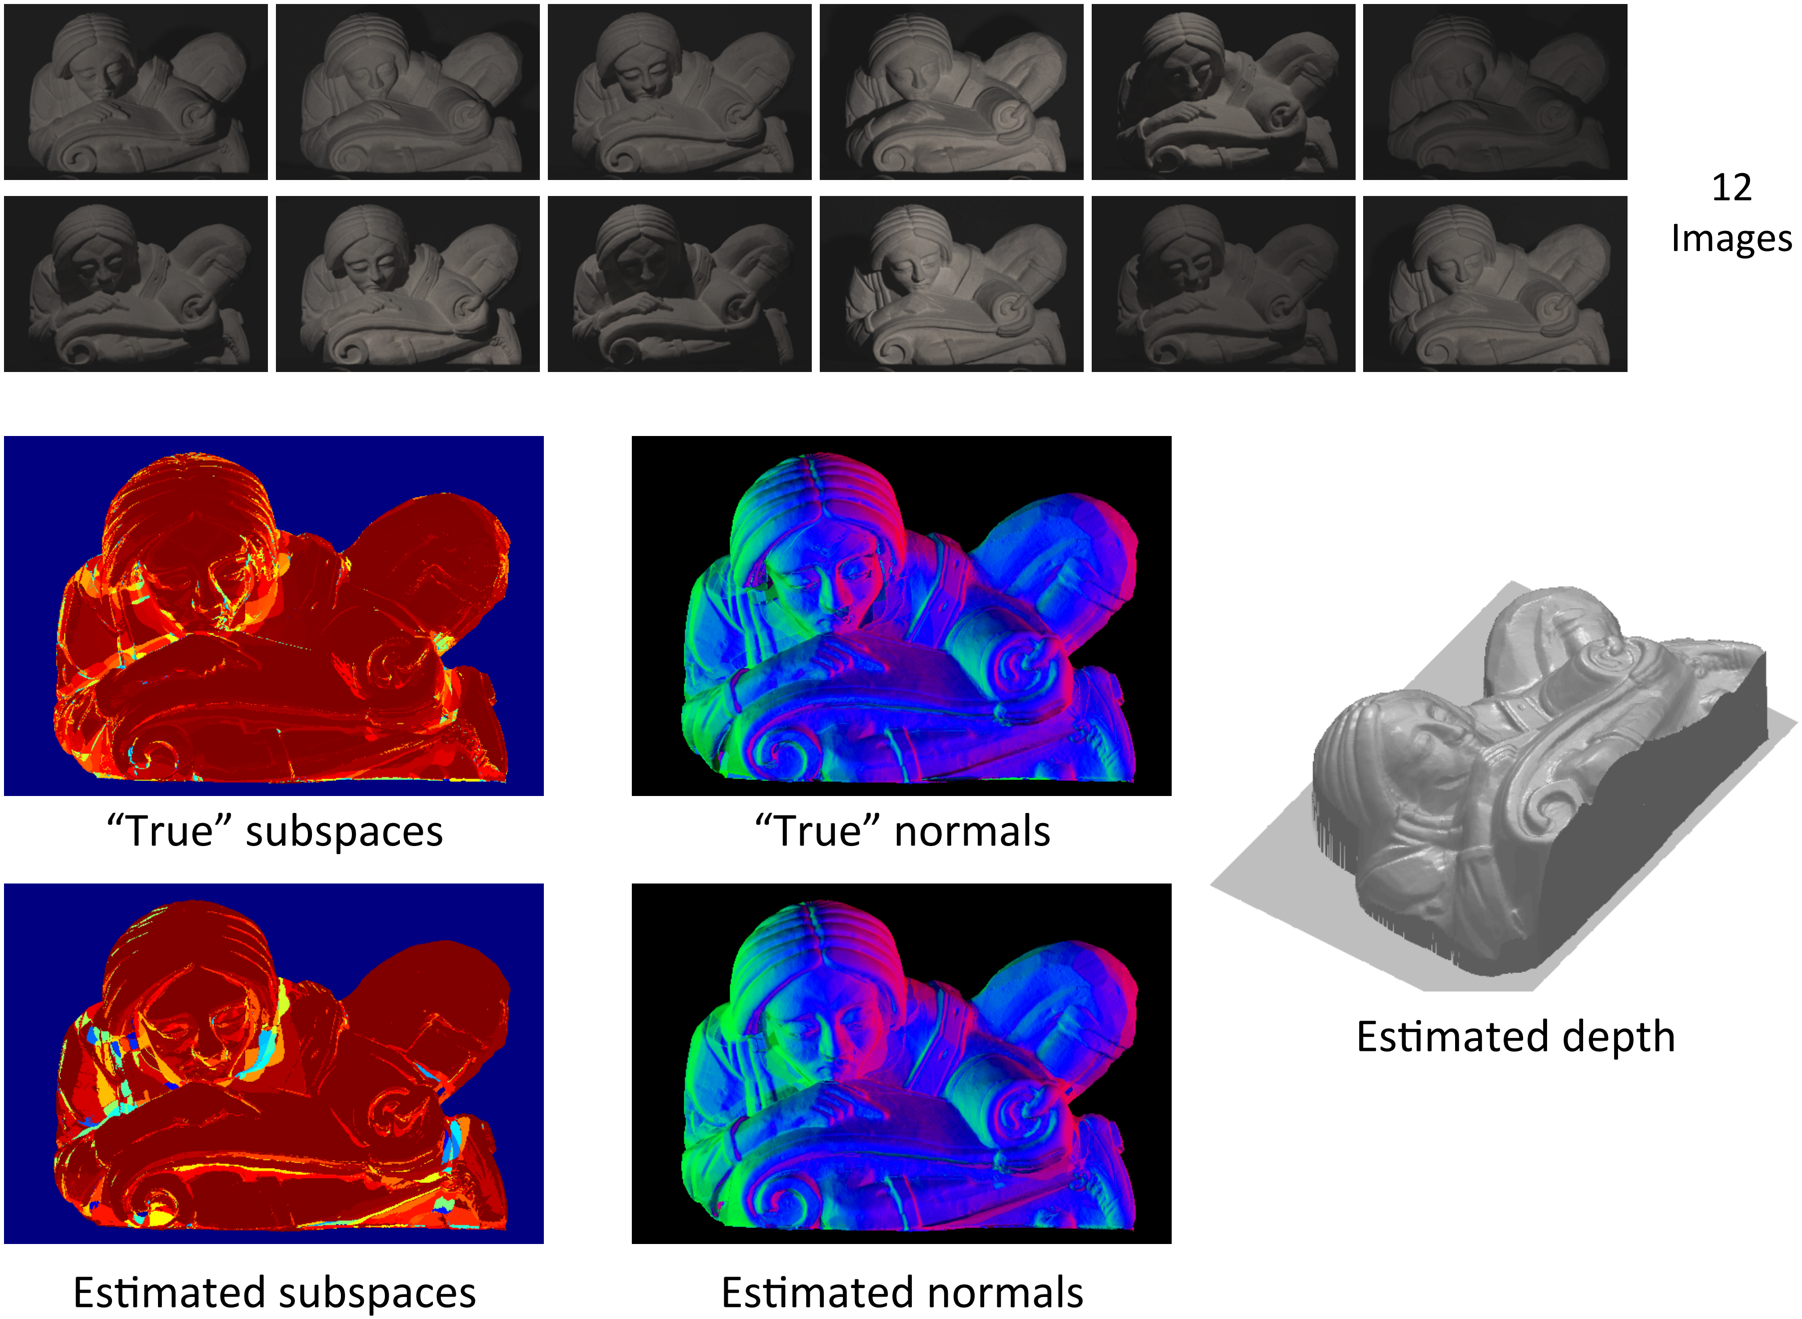
12
Images
"True" subspaces
"True" normals
Estimated depth
Estimated subspaces
Estimated normals

Some issues

- Degeneracies
  - Rank-deficient normals
  - Explicitly handle these in subspace estimation and normal recovery
- Deviations from Lambertian reflectance
  - Specify RANSAC error threshold appropriately
- Stability of subspace estimation
  - Intersections between subspaces
  - Large number of images, complex geometry

Conclusions

- An analysis of the influence of shadows on scene appearance.
- A novel bound on the dimensionality of scene appearance in the presence of shadows.
- An uncalibrated Photometric Stereo algorithm that is robust to shadowing.

- Extend analysis to mutual illumination
- Add spatial constraints
- Extend to more general cases (arbitrary BRDFs and illumination)

Thank you!

http://gvi.seas.harvard.edu