Illumination Cones and Uncalibrated Photometric Stereo

Computer Vision I
CSE252A
Lecture 8

What is the set of images of an object under all possible lighting conditions?

In answering this question, we'll arrive at a method for reconstructing surface shape w/ unknown lighting.

The Space of Images

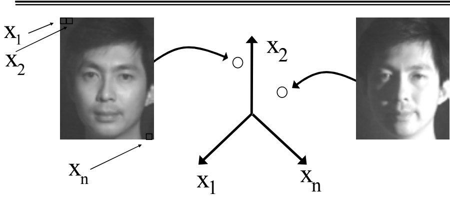

- Consider an n-pixel image to be a point in an n-dimensional space, $\mathbf{x} \in \mathbf{R}^n$.
- Each pixel value is a coordinate of **x**.
- Many results will apply to linear transformations of image space (e.g. filtered images)
- Other image representations (e.g. Cayley-Klein spaces, See Koenderink's "pixel f#@king paper")

Assumptions

For discussion, we assume:

- Lambertian reflectance functions.
- Objects have convex shape.
- Light sources at infinity.
- Orthographic projection.

- Note: many of these can be relaxed….

Lambertian Surface

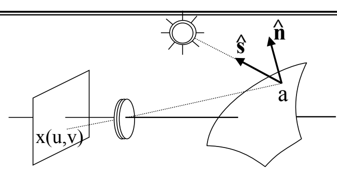

At image location (u,v), the intensity of a pixel x(u,v) is:

$$x(u,v) = [a(u,v)\, \hat{\mathbf{n}}(u,v)] \bullet [s_0 \hat{\mathbf{s}}]$$
$$= \mathbf{b}(u,v) \bullet \mathbf{s}$$

where

- a(u,v) is the albedo of the surface projecting to (u,v).
- **n**(u,v) is the direction of the surface normal.
- s_0 is the light source intensity.
- **s** is the direction to the light source.

Model for Image Formation

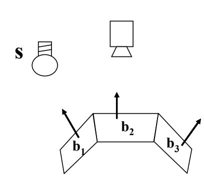

Lambertian Assumption with shadowing:

$$\mathbf{x} = \max(\mathbf{B}\,\mathbf{s}, 0) \qquad \mathbf{B} = \begin{bmatrix} -\ \mathbf{b}^T_1\ - \\ -\ \mathbf{b}^T_2\ - \\ \dots\dots \\ -\ \mathbf{b}^T_n\ - \end{bmatrix}_{n \times 3}$$

where

- **x** is an *n*-pixel image vector
- **B** is a matrix whose rows are unit normals scaled by the albedos
- $\mathbf{s} \in \mathbf{R}^3$ is vector of the light source direction scaled by intensity

3-D Linear subspace

The set of images of a Lambertian surface with no shadowing is a subset of 3-D linear subspace.

[Moses 93], [Nayar, Murase 96], [Shashua 97]

$L = \{\mathbf{x} \mid \mathbf{x} = B\mathbf{s}, \forall \mathbf{s} \in R^3\}$

where B is a n by 3 matrix whose rows are product of the surface normal and Lambertian albedo

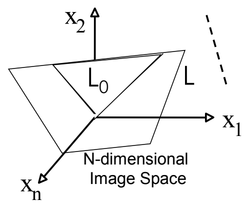

How do you construct subspace?

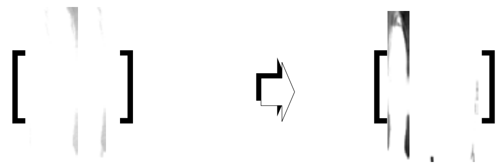

[$\mathbf{x}_1$ $\mathbf{x}_2$ $\mathbf{x}_3$] **B**

- Any three images w/o shadows taken under different lighting span L
- Not orthogonal
- Orthogonalize with Grahm-Schmidt

How do you construct subspace?

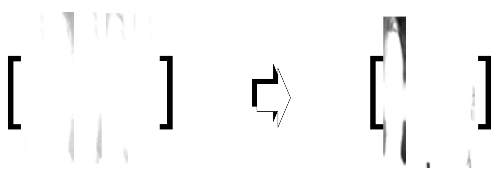

[$\mathbf{x}_1$ $\mathbf{x}_2$ $\mathbf{x}_3$ $\mathbf{x}_4$ $\mathbf{x}_5$] **B**

With more than three images, perform least squares estimation of **B** using Singular Value Decomposition (SVD)

Matrix Decompositions

- Definition: The factorization of a matrix ***M*** into two or more matrices $\boldsymbol{M}_1$, $\boldsymbol{M}_2$,…, $\boldsymbol{M}_n$, such that $\boldsymbol{M} = \boldsymbol{M}_1\boldsymbol{M}_2\ldots\boldsymbol{M}_n$.
- Many decompositions exist…
  - ***QR*** Decomposition
  - ***LU*** Decomposition
  - ***LDU*** Decomposition
  - Etc.

Singular Value Decomposition

Excellent ref: 'Matrix Computations," Golub, Van Loan

- Any *m* by *n* matrix ***A*** may be factored such that

$$\boldsymbol{A} = \boldsymbol{U}\Sigma\boldsymbol{V}^T$$

$$[m \times n] = [m \times m][m \times n][n \times n]$$

- ***U***: *m* by *m*, orthogonal matrix
  - Columns of **U** are the eigenvectors of $\boldsymbol{AA}^T$
- ***V***: *n* by *n*, orthogonal matrix,
  - columns are the eigenvectors of $\boldsymbol{A}^T\boldsymbol{A}$
- **Σ:** *m* by *n*, diagonal with non-negative entries (σ_1, σ_2, …, σ_s) with s=min(m,n) are called the called the singular values
  - Singular values are the square roots of eigenvalues of both $\boldsymbol{AA}^T$ and $\boldsymbol{A}^T\boldsymbol{A}$
  - *Result of SVD algorithm:* $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_s$

SVD Properties

- In Matlab [u s v] = svd(A), and you can verify that: A=u*s*v'
- r=Rank(***A***) = # of non-zero singular values.
- ***U***, ***V*** give us orthonormal bases for the subspaces of ***A***:
  - 1st *r* columns of ***U***: Column space of ***A***
  - Last *m - r* columns of ***U***: Left nullspace of ***A***
  - 1st *r* columns of ***V***: Row space of ***A***
  - last *n - r* columns of ***V***: Nullspace of ***A***
- *For d<= r, the first d column of U provide the best d-dimensional basis for columns of A in least squares sense.*

Thin SVD

- Any m by n matrix $\boldsymbol{A}$ may be factored such that

$$\boldsymbol{A} = \boldsymbol{U}\Sigma\boldsymbol{V}^T$$

$$[m \times n] = [m \times n][n \times n][n \times n]$$

- If m>n, then one can view Σ as:

$$\begin{bmatrix} \Sigma' \\ 0 \end{bmatrix}$$

- Where $\Sigma'=\text{diag}(\sigma_1, \sigma_2, \ldots, \sigma_s)$ with s=min(m,n), and lower matrix is (n-m by m) of zeros.
- Alternatively, you can write:

$$\boldsymbol{A} = \boldsymbol{U}'\Sigma'\boldsymbol{V}^T$$

- In Matlab, thin SVD is:[U S V] = svds(A)

Application: Pseudoinverse

- Given $\boldsymbol{y} = \boldsymbol{A}\mathbf{x}$, $\mathbf{x} = \boldsymbol{A}^+\mathbf{y}$
- For square $\boldsymbol{A}$, $\boldsymbol{A}^+ = \boldsymbol{A}^{-1}$
- For any $\boldsymbol{A}$…

$$\boldsymbol{A}^+ = \boldsymbol{V}\Sigma^{-1}\boldsymbol{U}^T$$

- $\boldsymbol{A}^+$ is called the ***pseudoinverse*** of ***A.***
- $\mathbf{x} = \boldsymbol{A}^+\mathbf{y}$ is the least-squares solution of $\boldsymbol{y} = \boldsymbol{A}\mathbf{x}$.
- Alternative to previous solution.

Estimating B with SVD

1. Construct data matrix

$$\mathbf{D} = [\ \mathbf{x}_1\ \mathbf{x}_2\ \mathbf{x}_3 \ldots \mathbf{x}_n\ ]$$

2. [**U S V**] = svds(D)

- If data had no noise, then rank(D)=3, and the first three singular values (S) would be positive and rest would be zero.
- Take first three column of **u** as **B.**

Still Life

Original Images

Basis Images

Rendering Images: $\Sigma_i \max(\mathbf{B}\mathbf{s}_i, \mathbf{0})$

1 Light 2 Lights 3 Lights

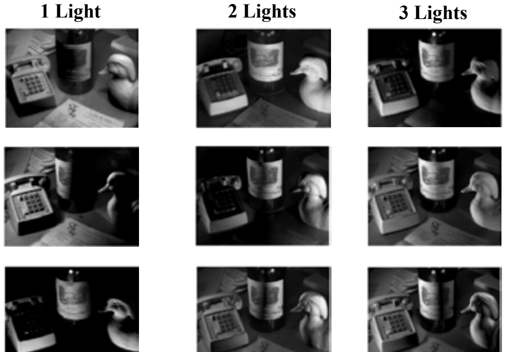

Face Basis

Original Images

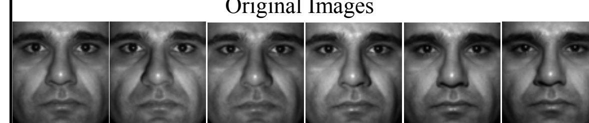

Basis Images

Movie with Attached Shadows

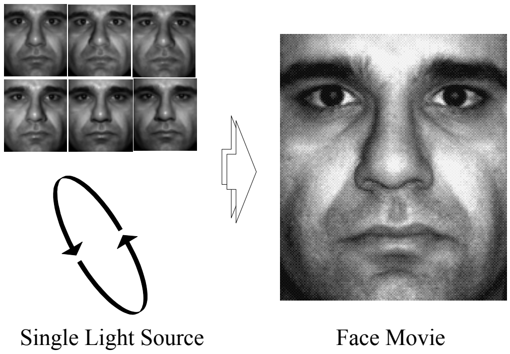

3-D Linear subspace

The set of images of a Lambertian surface with no shadowing is a subset of 3-D linear subspace.

[Moses 93], [Nayar, Murase 96], [Shashua 97]

$L = \{\mathbf{x} \mid \mathbf{x} = B\mathbf{s}, \forall \mathbf{s} \in R^3\}$

where B is a *n* by 3 matrix whose rows are product of the surface normal and Lambertian albedo

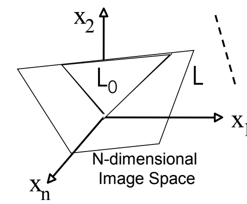

Set of Images from a Single Light Source

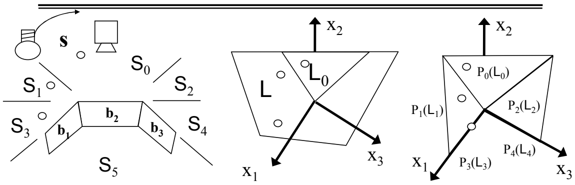

- Let L_i denote the intersection of L with an orthant *i* of R^n.
- Let P_i (L_i) be the projection of L_i onto a "wall" of the positive orthant given by max(**x**, **0**).

Then, the set of images of an object produced by a single light source is:

$$U = \bigcup_{i=0}^{M} P_i(L_i)$$

Lambertian, Shadows and Multiple Lights

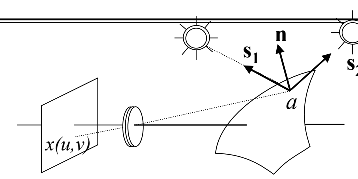

The image **x** produced by multiple light sources is

$$\mathbf{x} = \sum_i \max(\mathbf{B}\,\mathbf{s}_i, 0)$$

where

- **x** is an *n*-pixel image vector.
- **B** is a matrix whose rows are unit normals scaled by the albedo.
- $\mathbf{s}_i$ is the direction and strength of the light source *i*.

Set of Images from Multiple Light Sources

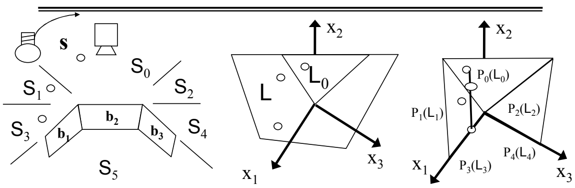

- With two lights on, resulting image along line segment between single source images: superposition of images, non-negative lighting
- For all numbers of sources, and strengths, rest is convex hull of U.

The Illumination Cone

Theorem:: *The set of images of any object in fixed posed, but under all lighting conditions, is a* **convex cone** *in the image space.*

(Belhumeur and Kriegman, IJCV, 98)

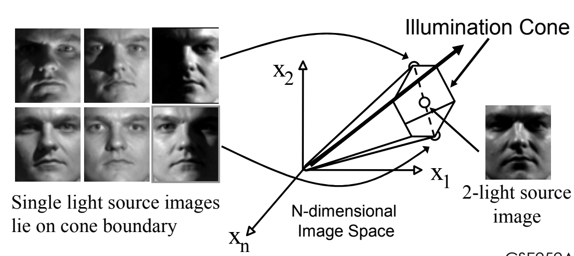

Some natural ideas & questions

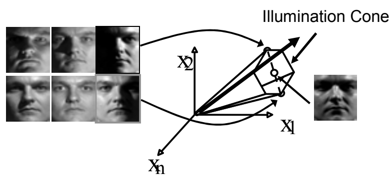

- Can the cones of two different objects intersect?
- Can two different objects have the same cone?
- How big is the cone?
- How can cone be used for recognition?

Do Ambiguities Exist?

Can two objects of differing shapes produce the same illumination cone?

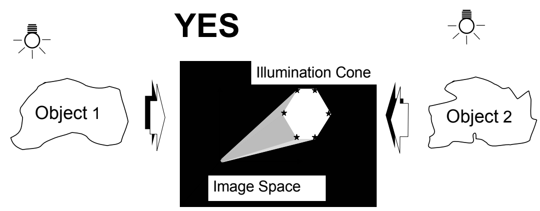

Do Ambiguities Exist? Yes

- Cone is determined by linear subspace **L**
- The columns of **B** span **L**
- For any **A** $\in$ GL(3), **B* = BA** also spans **L.**
- For any image of **B** produced with light source **S**, the same image can be produced by lighting **B*** with $\mathbf{S^*=A^{-1}S}$ because
  $\mathbf{X = B^*S^* = B\,AA^{-1}S = BS}$
- When we estimate B using SVD, the rows are NOT generally normal * albedo.
-

Surface Integrability

In general, $\mathbf{B}^*$ does not have a corresponding surface.

Linear transformations of the surface normals in general do not produce an integrable normal field.

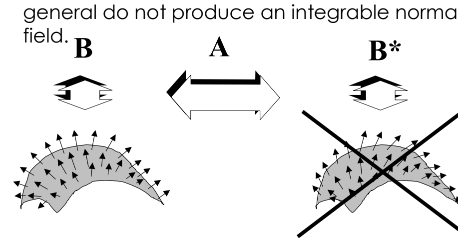

GBR Transformation

Only Generalized Bas-Relief transformations satisfy the integrability constraint:

$$\mathbf{A} = \mathbf{G}^T = \begin{bmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\nu \\ 0 & 0 & 1 \end{bmatrix}^T$$

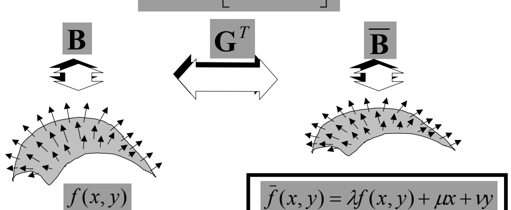

Generalized Bas-Relief Transformations

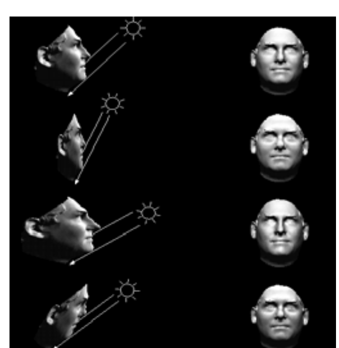

Objects differing by a GBR have the same illumination cone.

Without knowledge of light source location, one can only recover surfaces up to GBR transformations.

Uncalibrated photometric stereo

1. Take n images as input, perform SVD to compute B*.
2. Find some A such that B*A is close to integrable.
3. Integrate resulting gradient field to obtain height function f*(x,y).

Comments:

- f*(x,y) differs from f(x,y) by a GBR.
- Can use specularities to resolve GBR for non-Lambertian surface.

What about cast shadows for nonconvex objects?

P.P. Reubens in Opticorum Libri Sex, 1613

GBR Preserves Shadows

Given a surface *f* and a GBR transformed surface *f'* then for every light source **s** which illuminates *f* there exists a light source **s'** which illuminates *f'* such that the **attached** and **cast shadows** are identical.

GBR is the only transform that preserves shadows.

[Kriegman, Belhumeur 2001]

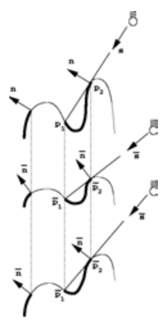

Bas-Relief Sculpture

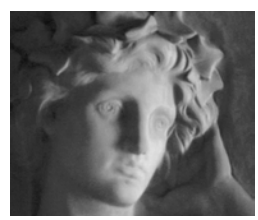

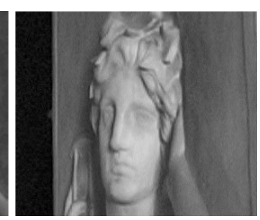

Codex Urbinas

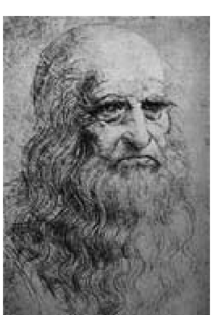

As far as light and shade are concerned low relief fails both as sculpture and as painting, because** the shadows correspond to the low nature of the relief, **as for example in the shadows of foreshortened objects, which will not exhibit the depth of those in painting or in sculpture in the round.

Leonardo da Vinci
Treatise on Painting (Kemp)