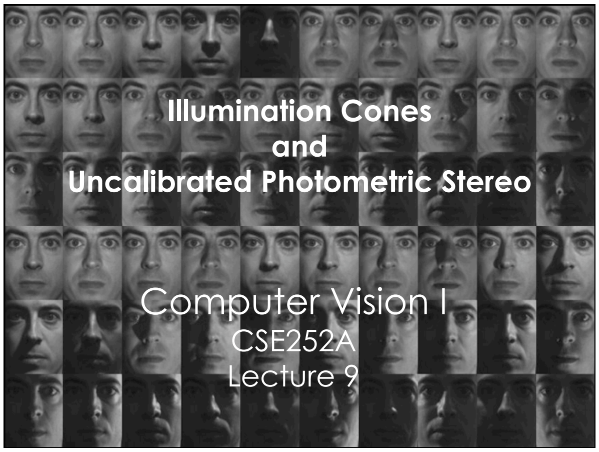

Illumination Cones and Uncalibrated Photometric Stereo

Computer Vision I
CSE252A
Lecture 9

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

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

CSE252A

The Space of Images

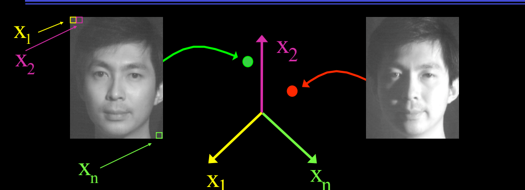

- Consider an n-pixel image to be a point in an n-dimensional space, $\mathbf{x} \in \mathbb{R}^n$.
- Each pixel value is a coordinate of $\mathbf{x}$.
- Many results will apply to linear transformations of image space (e.g. filtered images)
- Other image representations (e.g. Cayley-Klein spaces, See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

CSE252A

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....

CSE252A

Lambertian Surface

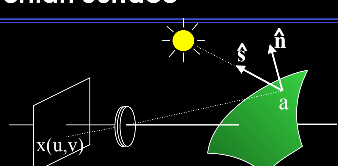

At image location (u,v), the intensity of a pixel x(u,v) is:

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

where

- $a(u,v)$ is the albedo of the surface projecting to (u,v).
- $\mathbf{n}(u,v)$ is the direction of the surface normal.
- s_0 is the light source intensity.
- $\mathbf{s}$ is the direction to the light source.

CSE252A

Model for Image Formation

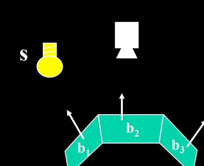

Lambertian Assumption with shadowing:

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

where

- $\mathbf{x}$ is an n-pixel image vector
- $\mathbf{B}$ 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

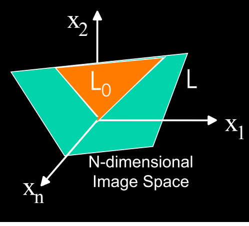

How do you construct subspace?

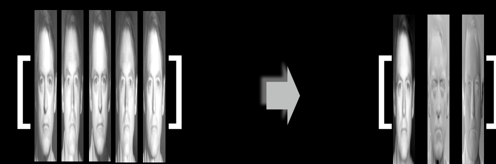

$[\, x_1\ x_2\ x_3\ x_4\ x_5 \,]$ **B**

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

Still Life

Original Images

Basis Images

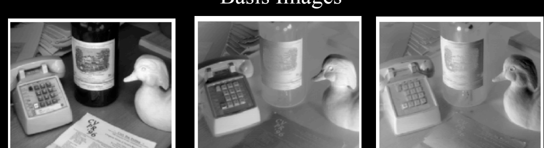

Rendering Images: $\Sigma_i \max(\mathbf{B}\mathbf{s}_i, \mathbf{0})$

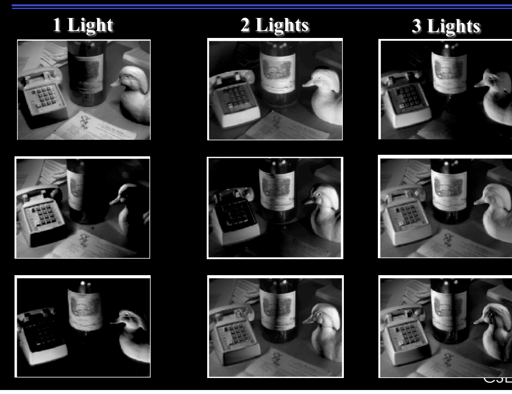

3-D Linear subspace

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

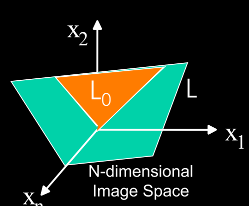

Set of Images from a Single Light Source

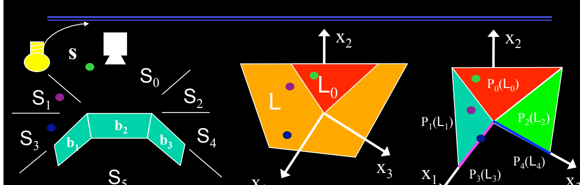

- Let L_i denote the intersection of L with an orthant *i* of R^n.
- Let $P_i(L_i)$ be the projection of L_i onto a "wall" of the positive orthant given by $\max(\mathbf{x}, \mathbf{0})$.

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

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

Lambertian, Shadows and Multiple Lights

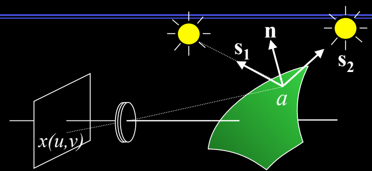

The image **x** produced by multiple light sources is

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

- **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*.

CSE252A

Set of Images from Multiple Light Sources

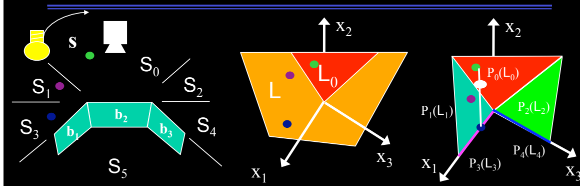

- With two lights on, resulting image along line segment between single source images: superposition of images, non-negative lighting
- For all numbers of sources, and strengths, rest is convex hull of U.

CSE252A

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)

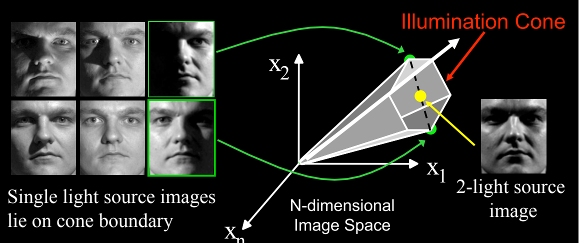

CSE252A

Some natural ideas & questions

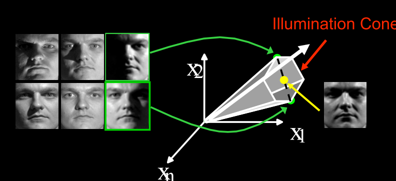

- Can the cones of two different objects intersect?
- Can two different objects have the same cone?
- How big is the cone?
- How can cone be used for recognition?

CSE252A

Do Ambiguities Exist?

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

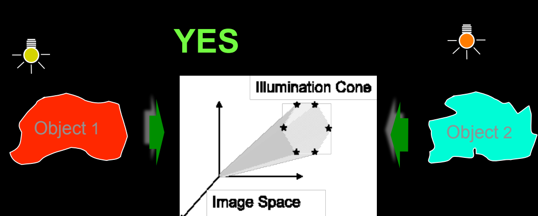

CSE252A

Do Ambiguities Exist? Yes

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

CSE252A

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.

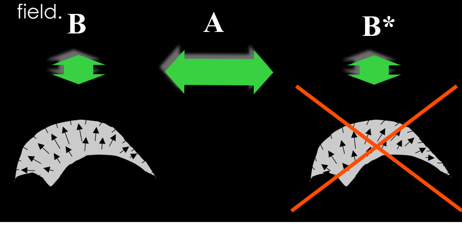

CSE252A

GBR Transformation

Only Generalized Bas-Relief transformations satisfy the integrability constraint:

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

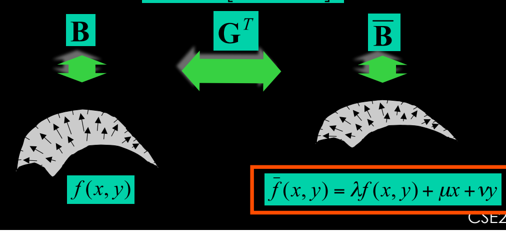

CSE252A

Generalized Bas-Relief Transformations

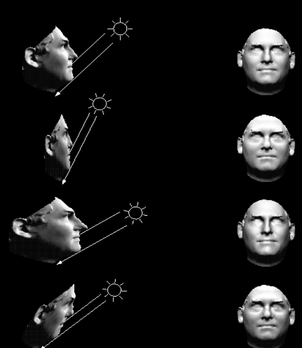

Objects differing by a GBR have the same illumination cone.

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

CSE252A

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.

CSE252A

What about cast shadows for nonconvex objects?

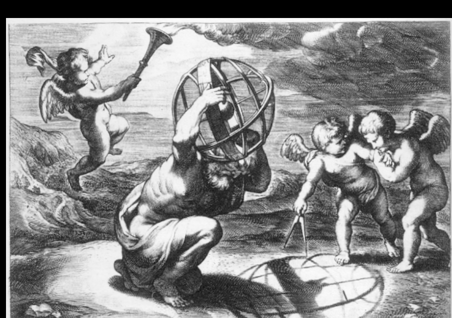

P.P. Reubens in Opticorum Libri Sex, 1613

CSE252A

GBR Preserves Shadows

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

GBR is the only transform that preserves shadows.

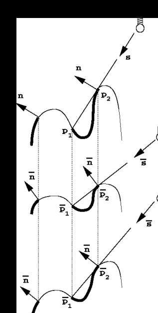

[Kriegman, Belhumeur 2001]

CSE252A

Bas-Relief Sculpture

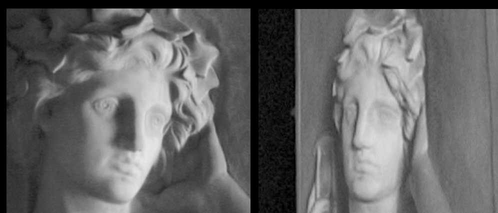

CSE252A

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)

CSE252A

The Illumination Cone

Thm: The span of the extreme rays of the illumination cone is equal to the number of distinct surface normals – i.e., as high as n.

The number of extreme rays of the cone is n(n-1)+2

(Belhumeur and Kriegman, IJCV,'98)

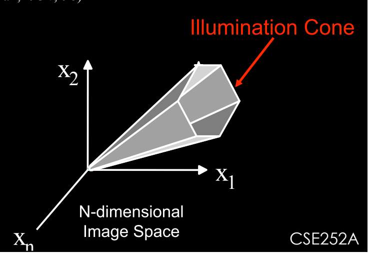

CSE252A

Shape of the Illumination Cone

Observation: The illumination cone is **flat** with most of its volume concentrated near a **low-dimensional** linear subspace

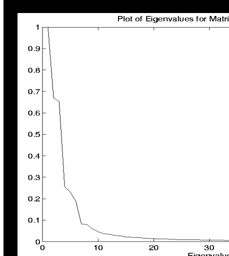

	Ball	Face	Phone	Parrot
#1	48.2	53.7	67.9	42.8
#3	94.4	90.2	88.2	76.3
#5	97.9	93.5	94.1	84.7
#7	99.1	95.3	96.3	88.5
#9	99.5	96.3	97.2	90.7

[Epstein, Hallinan, Yuille 95]

Dimension:

CSE252A

Recent results

- Illumination cone is well capture by nine dimensions for a convex Lambertian surface.
  - Spherical Harmonic representation of lighting & BRDF.

CSE252A

Illumination Cones: Recognition Method

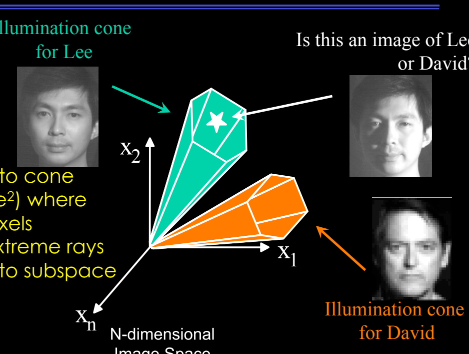

- Distance to cone
- Cost O(ne^2) where
  - n: # pixels
  - e: # extreme rays
- Distance to subspace

CSE252A

Generating the Illumination Cone

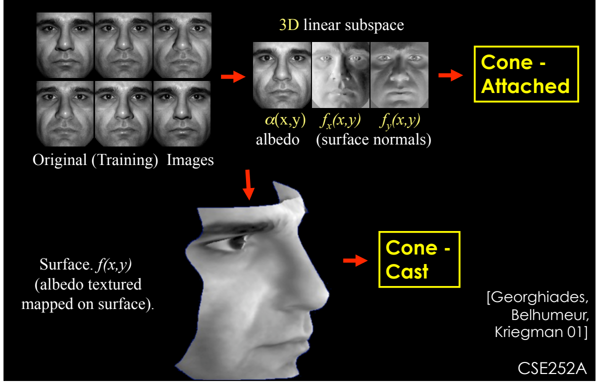

[Georghiades, Belhumeur, Kriegman 01]

CSE252A

Predicting Lighting Variation

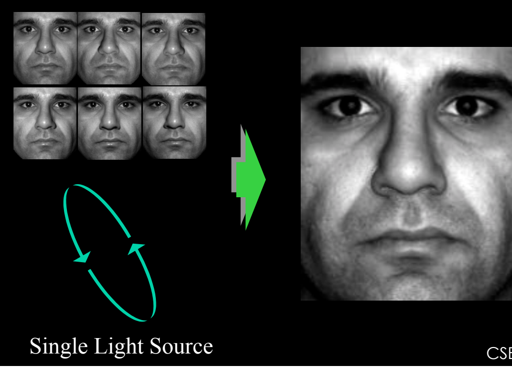

CSE252A

Yale Face Database B

64 Lighting Conditions
9 Poses
=> 576 Images per Person

CSE252A

Face Recognition: Test Subsets

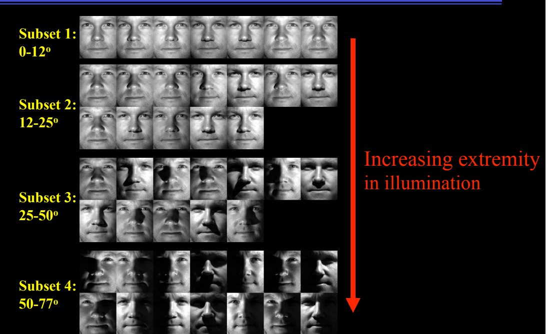

Test images divided into 4 subsets depending on illumination.

CSE252A

Geodesic Dome Database - Frontal Pose

[Georghiades, Belhumeur, Kriegman 01]

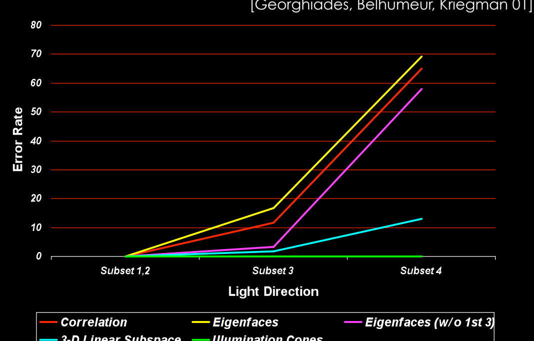

CSE252A

Face Recognition: Lighting & Pose

1. Union of linear subspaces
   1. Sample the pose space, and for each pose construct illumination cone.
   2. Cone can be approximated by linear subspace (used 11-D)
2. For computational efficiency, project using PCA to 100-D

CSE252A

Illumination Variability Reveals Shape

CSE252A

Pose Variability in test set: Up to 24°

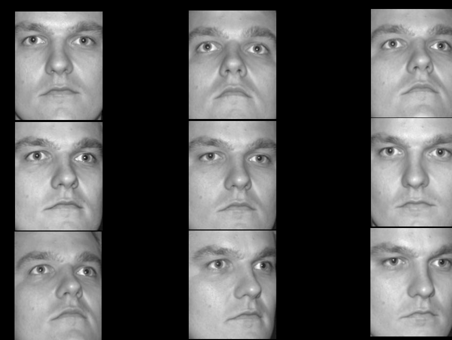

CSE252A

Illumination Cone Face Recognition Result: Pose and Lighting

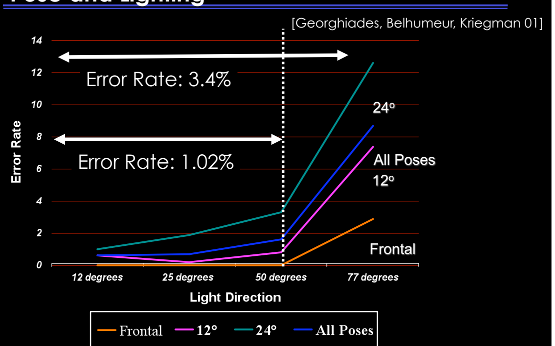

CSE252A

Closest Sample to Test Image: Examples

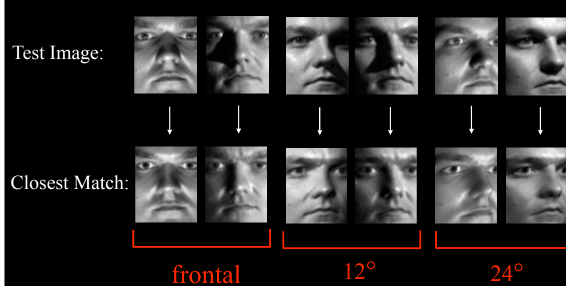

CSE252A

Illumination & Image Set

- Lack of illumination invariants
  [Chen, Jacobs, Belhumeur 98]
- Set of images of Lambertian surface w/o shadowing is 3-D linear subspace
  [Moses 93], [Nayar, Murase 96], [Shashua 97]
- Empirical evidence that set of images of object is well-approximated by a low-dimensional linear subspace
  [Hallinan 94], [Epstein, Hallinan, Yuille 95]
- Illumination cones
  [Belhumeur, Kriegman 98]
- Spherical harmonics lighting & images
  [Basri, Jacobs 01], [Ramamoorthi, Hanrahan 01]
- Analytic PCA of image over lighting
  [Ramamoorthi 02]

CSE252A

Some subsequent work

1. "Face Recognition Under Variable Lighting using Harmonic Image Exemplars," Zhang, Samaras, CVPR03
2. "Clustering Appearances of Objects Under Varying Illumination Conditions," Ho, Lee, Lim, Kriegman, CVPR 03
3. "Low-Dimensional Representations of Shaded Surfaces under Varying Illumination, Nillius, Eklundh, CVPR03
4. "Using Specularities for Recognition," Osadchy, Jacobs, Ramamoorthi, ICCV 03

CSE252A