Resolving the Generalized Bas-Relief Ambiguity by Entropy Minimization

Neil G. Alldrin Satya P. Mallick David J. Kriegman
University of California, San Diego

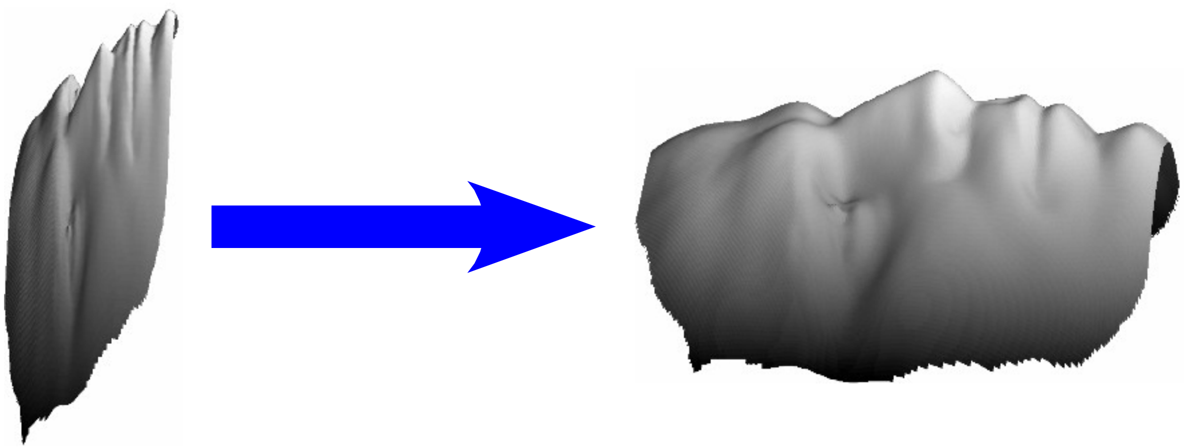

UCSD Jacobs | Computer Science and Engineering

Image Formation Model

- Camera centered coordinate system
- Orthographic camera

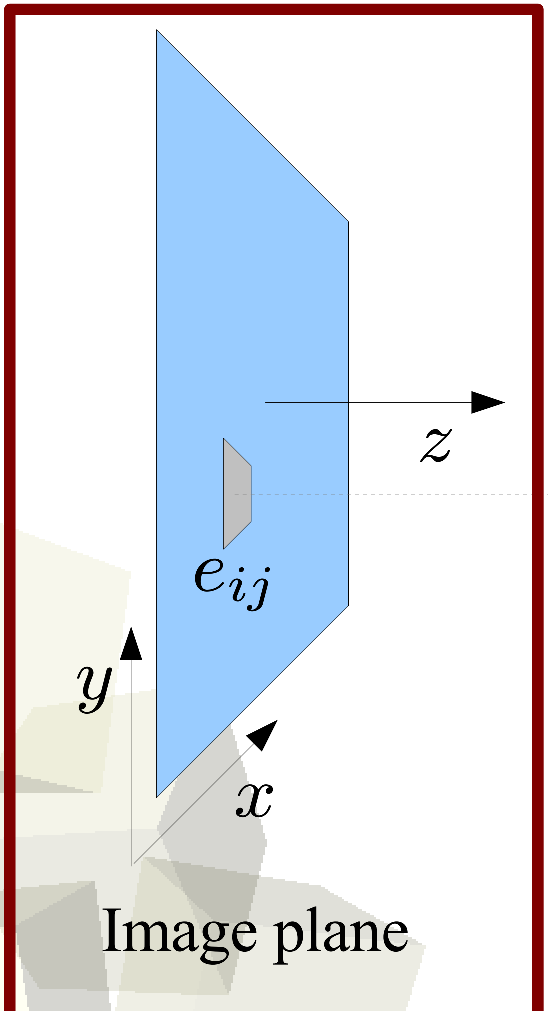

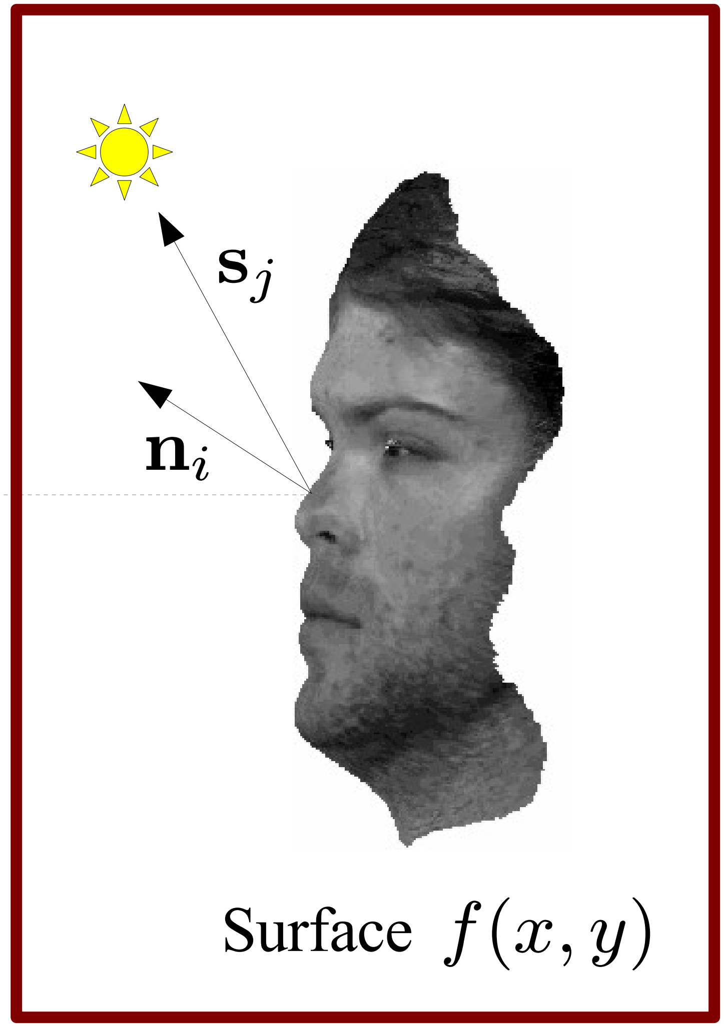

Image Formation Model

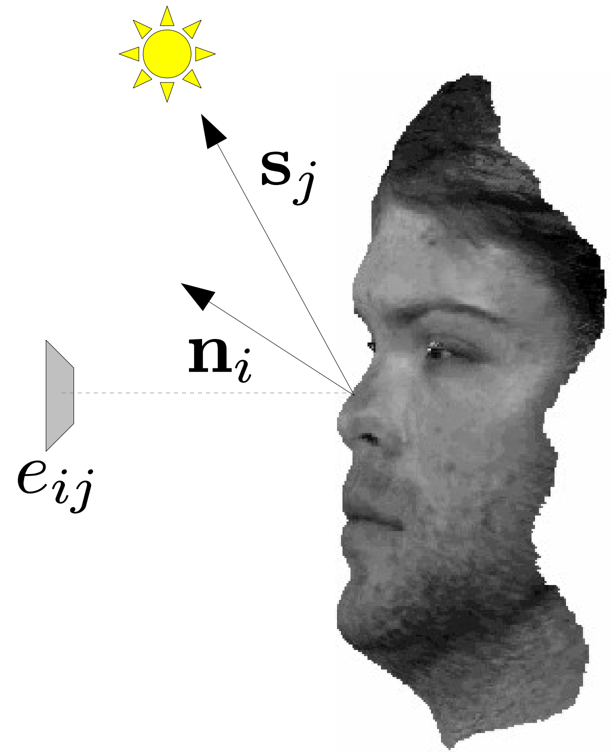

- Assumptions :
  - Lambertian reflectance
  - Orthographic camera
  - Distant point light source
  - No cast shadows / interreflections
- Image formation,
  - $e_{ij} = \rho_i \mathbf{n}_i^\top \mathbf{s}_j = \mathbf{b}_i^\top \mathbf{s}_j$
  - e_{ij} – pixel intensity at ith pixel, jth image
    $\mathbf{n}_i$ – surface normal at ith pixel (3×1)
    ρ_i – albedo at ith pixel
    $\mathbf{b}_i$ – facet vector $\mathbf{b}_i = \rho_i \mathbf{n}_i$ (3×1)
    $\mathbf{s}_j$ – light source vector of jth image (3×1)

UCSD Jacobs | Computer Science and Engineering

Photometric Stereo

- [Silver 1980, Woodham 1981]
- Goal :
  - Recover surface (normal map)
  - Fixed scene
  - Fixed viewpoint
  - Varying Illumination
- Solve linear system,
  - $\mathbf{E} = \mathbf{B}\mathbf{S}$
  - $\{\mathbf{E}\}_{i,j} = e_{ij}$
    $\{\mathbf{B}\}_{i,:} = \rho_i \mathbf{n}_i^\top$
    $\{\mathbf{S}\}_{:,j} = \mathbf{s}_j$
- Solution : $\mathbf{B} = \mathbf{E}\mathbf{S}^\dagger$

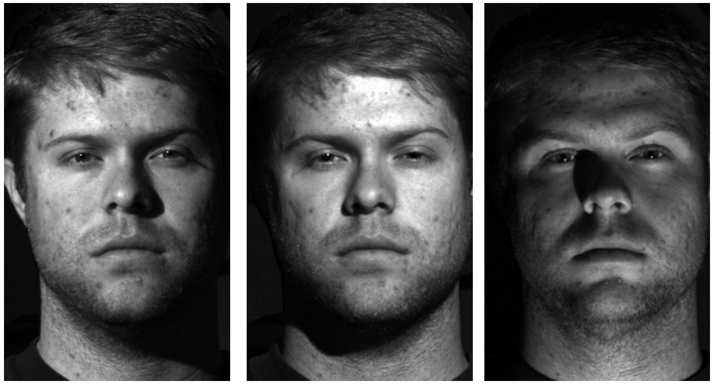

(Yale Face Database B)

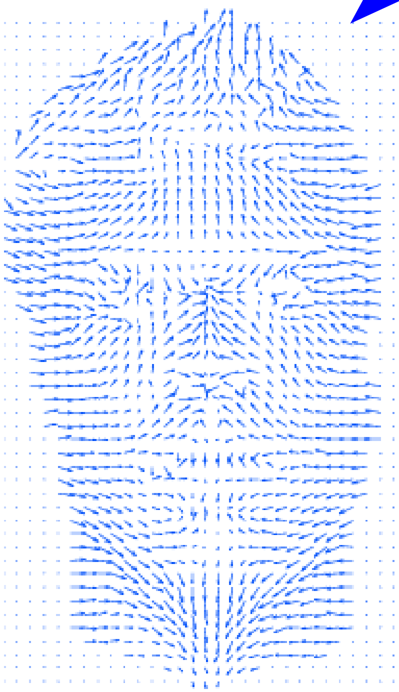

Uncalibrated Photometric Stereo

- What if the lighting $\mathbf{S}$ is unknown?

- Family of solutions,
  - $\mathbf{E} = \mathbf{B}\mathbf{S} = \mathbf{B}\mathbf{A}^{-1}\mathbf{A}\mathbf{S}$
  - $\mathbf{A} \in GL(3)$

    [Hayakawa 1994]
    [Epstein, Yuille, & Belhumeur 1996]
    [Rosenholtz & Koenderink 1996]

- Factorize $\mathbf{E}$ with rank 3 SVD approximation,
  - $\mathbf{E} = \mathbf{U}\mathbf{\Sigma}_{3\times 3}\mathbf{V}^{\top}$ ➡ $\mathbf{B}\mathbf{A}^{-1}\mathbf{A}\mathbf{S}$
  - Recovers $\mathbf{B}$, $\mathbf{S}$ up to a 3x3 invertible linear transform

The GBR Ambiguity

- SVD + Integrability,
  $$\mathbf{E} = \hat{\mathbf{B}}\hat{\mathbf{S}} = \mathbf{B}\mathbf{G}^{-1}\mathbf{G}\mathbf{S}$$
  [Belhumeur, Kriegman & Yuille 1997]
  [Yuille & Snow 1997]

- $\mathbf{G}$ encodes the GBR ambiguity,
  - $$\mathbf{G} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{pmatrix}$$
  - 3 parameters (μ, ν, λ)
  - $\mathbf{G} \subset GL(3)$

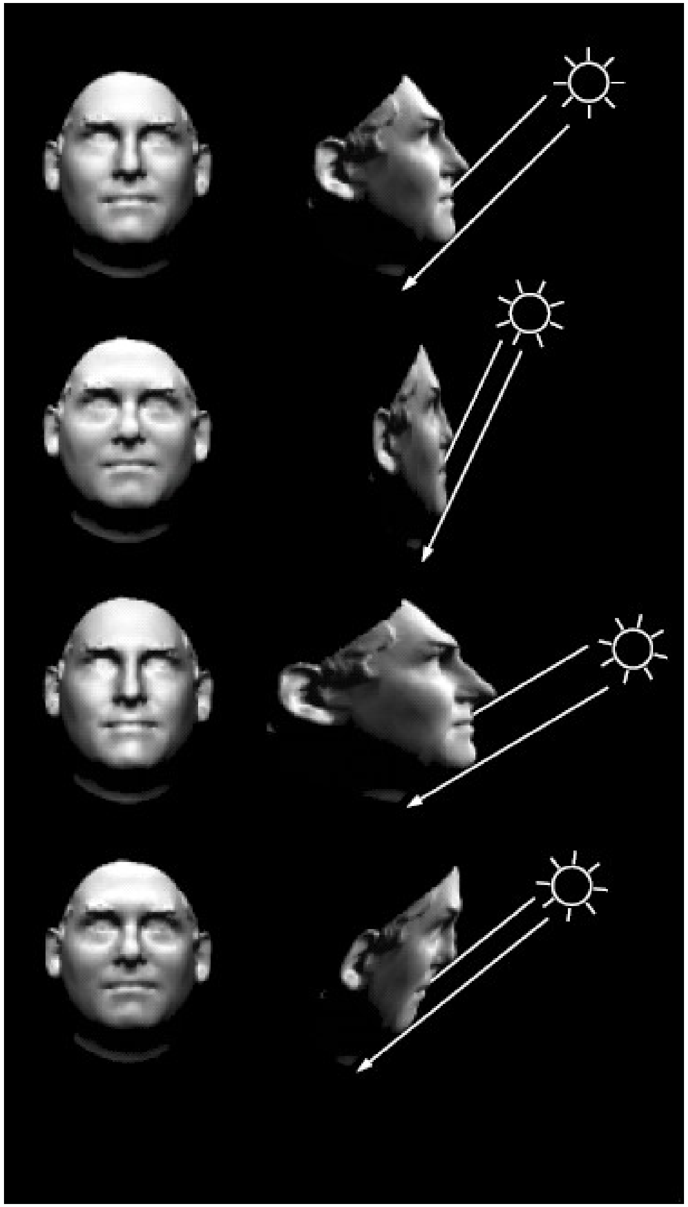

(Belhumeur, Kriegman, and Yuille)

UCSD Jacobs | Computer Science and Engineering

The GBR Ambiguity

- SVD + Integrability,

$$\mathbf{E} = \hat{\mathbf{B}}\hat{\mathbf{S}} = \mathbf{B}\mathbf{G}^{-1}\mathbf{G}\mathbf{S}$$

[Belhumeur, Kriegman & Yuille 1997]
[Yuille & Snow 1997]

- Surface height,

$$z = f(x, y)$$

$$\hat{z} = \lambda f(x, y) + \mu x + \nu y$$

- Albedo & normal,

$$\hat{\rho} = \rho \|\mathbf{n}^\top \mathbf{G}^{-1}\|$$

$$\hat{\mathbf{n}}^\top = \frac{\mathbf{n}^\top \mathbf{G}^{-1}}{\|\mathbf{n}^\top \mathbf{G}^{-1}\|}$$

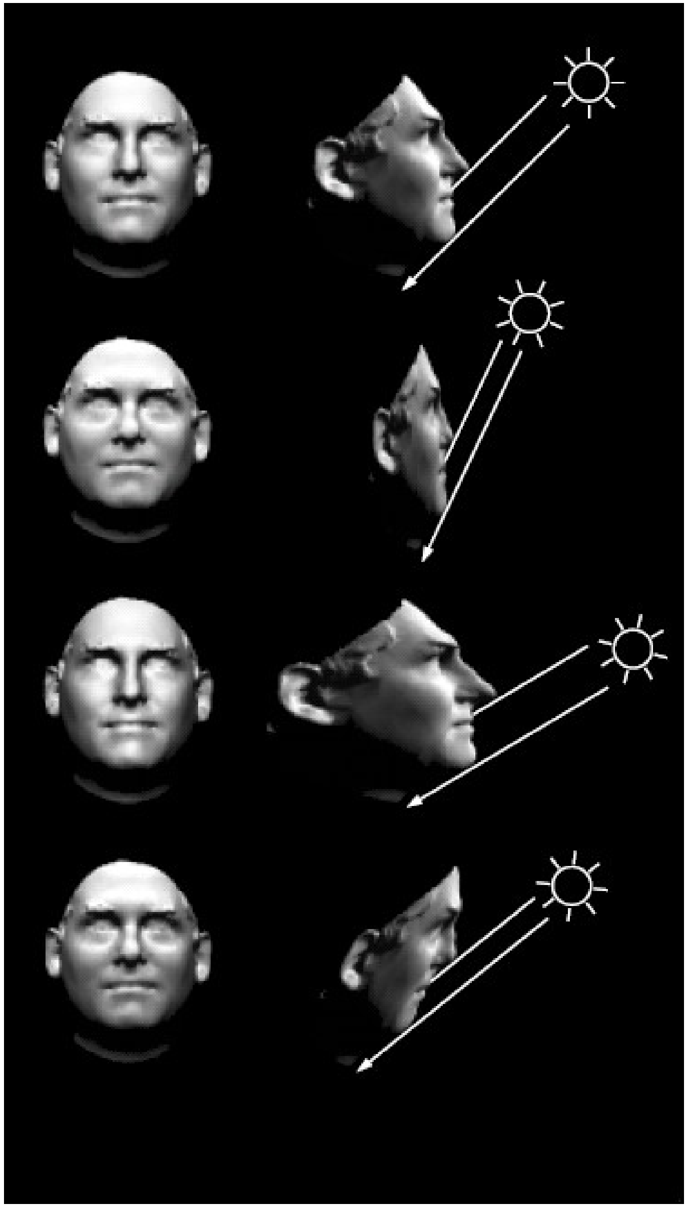

(Belhumeur, Kriegman, and Yuille)

UCSD Jacobs | Computer Science and Engineering

The GBR Ambiguity

- SVD + Integrability,

$$\mathbf{E} = \hat{\mathbf{B}}\hat{\mathbf{S}} = \mathbf{B}\mathbf{G}^{-1}\mathbf{G}\mathbf{S}$$

[Belhumeur, Kriegman & Yuille 1997]
[Yuille & Snow 1997]

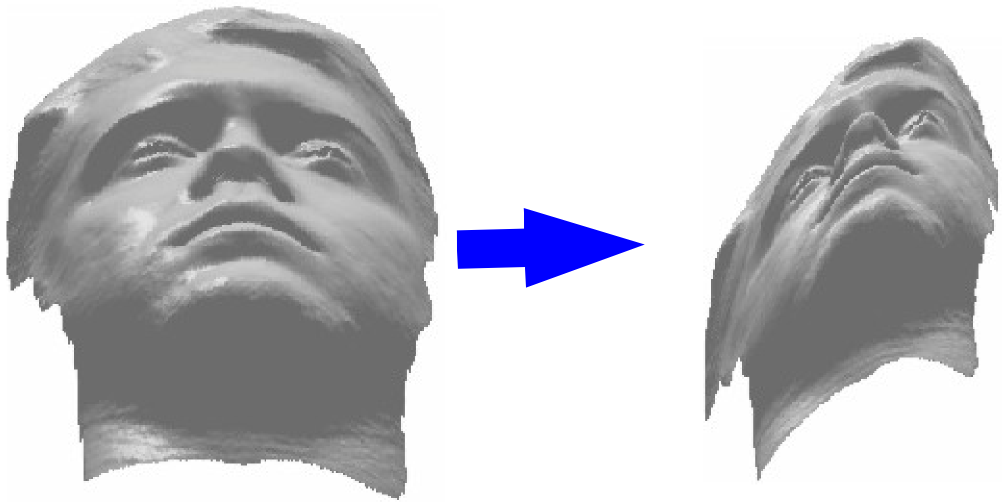

- Surface height,

$$z = f(x, y)$$

$$\hat{z} = \lambda f(x, y) + \mu x + \nu y$$

- Albedo & normal,

$$\hat{\rho} = \rho \|\mathbf{n}^\top \mathbf{G}^{-1}\|$$

$$\hat{\mathbf{n}}^\top = \frac{\mathbf{n}^\top \mathbf{G}^{-1}}{\|\mathbf{n}^\top \mathbf{G}^{-1}\|}$$

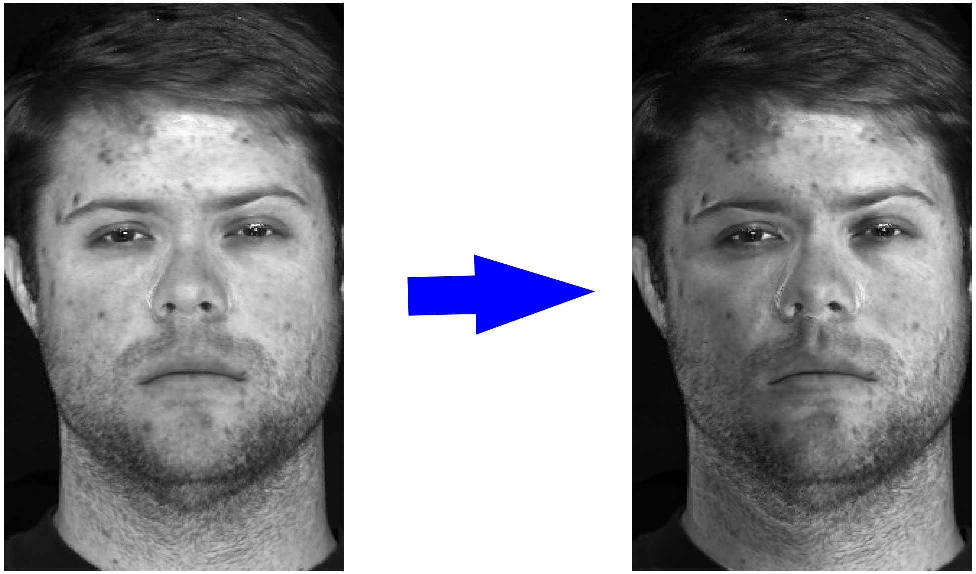

Albedo map — Albedo map (GBR)

The GBR Ambiguity

- SVD + Integrability,

$$\mathbf{E} = \hat{\mathbf{B}}\hat{\mathbf{S}} = \mathbf{B}\mathbf{G}^{-1}\mathbf{G}\mathbf{S}$$

[Belhumeur, Kriegman & Yuille 1997]
[Yuille & Snow 1997]

- Surface height,

$$z = f(x, y)$$

$$\hat{z} = \lambda f(x, y) + \mu x + \nu y$$

- Albedo & normal,

$$\hat{\rho} = \rho \|\mathbf{n}^\top \mathbf{G}^{-1}\|$$

$$\hat{\mathbf{n}}^\top = \frac{\mathbf{n}^\top \mathbf{G}^{-1}}{\|\mathbf{n}^\top \mathbf{G}^{-1}\|}$$

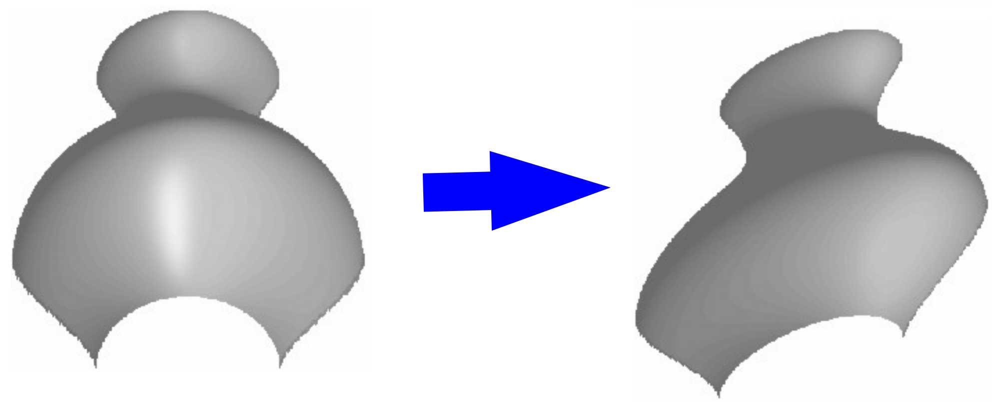

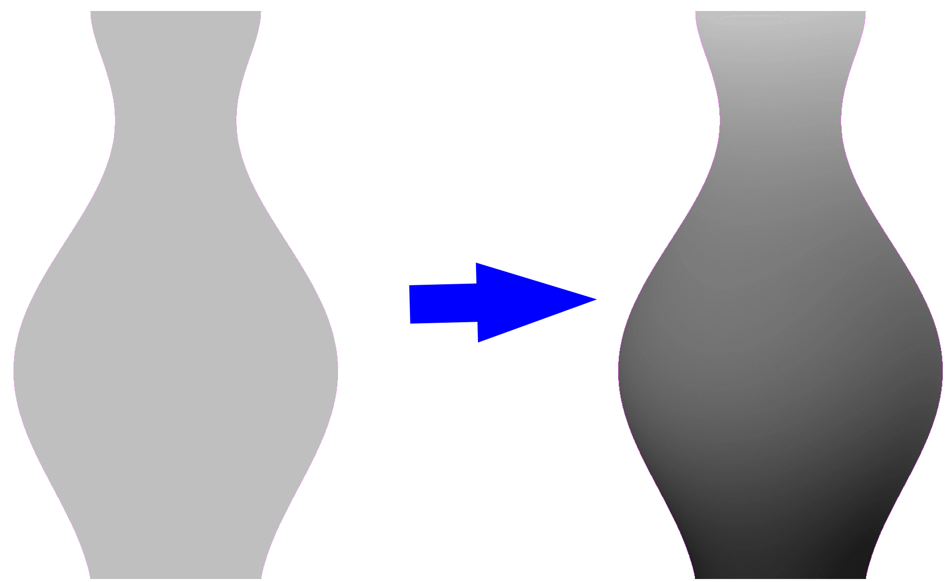

Albedo map

Albedo map (GBR)

Resolving the GBR

- Need additional constraints,
  - Light source strength
    [Yuille & Snow 1997]
  - Surface reflectance
    [Drbohlav & Sara 2002; Georghiades 2003; Tan et al. 2007]
  - Surface geometry
    [Georghiades et al. 2001]
  - Interreflections
    [Chandraker et al. 2005]
  - Albedo distribution
    [Hayakawa 1994]

- Albedo distribution :
  - Only uniform albedo has been exploited previously.

Motivation

- Consider an object with one albedo
- Distribution well approximated by a delta function

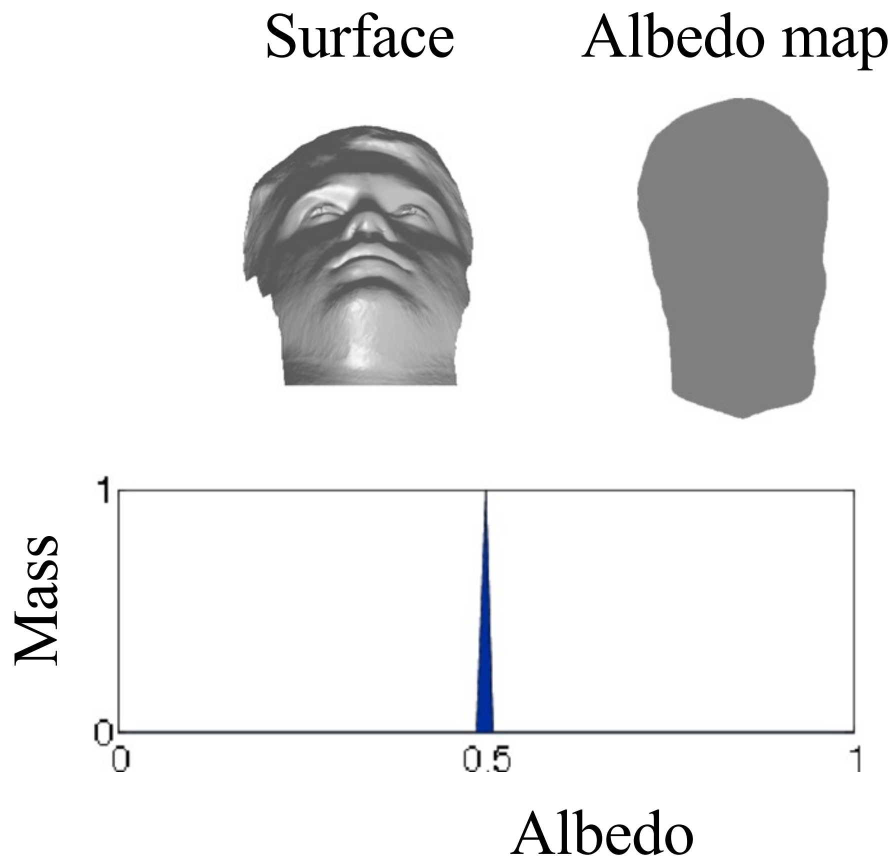

Motivation

- Under a GBR, the distribution is smeared
- Smearing dependent on $\mathbf{G}$ and the distribution of $\mathbf{n}$
  - $\hat{\rho} = \rho \|\mathbf{n}^\top \mathbf{G}^{-1}\|$
- Analytic solution :
  $\hat{\mathbf{b}}^\top \mathbf{G}\mathbf{G}^\top \hat{\mathbf{b}} = \rho^2$
- Linear in
  $(\mu, \nu, \mu^2 + \nu^2 + \lambda^2, \rho^2)$

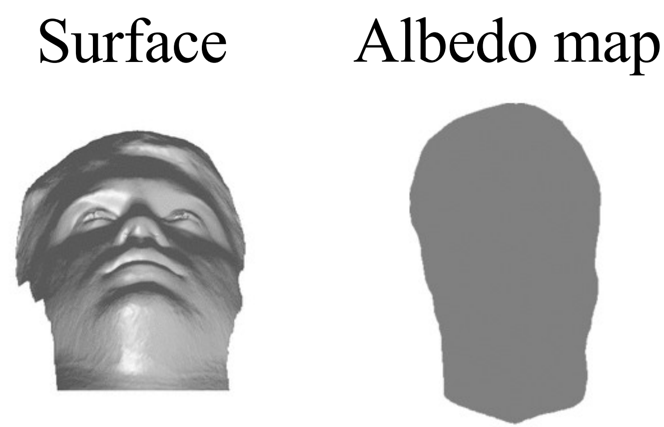

$$\begin{pmatrix} \mu : 0 \rightarrow 0.5 \\ \nu : 0 \\ \lambda : 1 \rightarrow 0.75 \end{pmatrix}$$

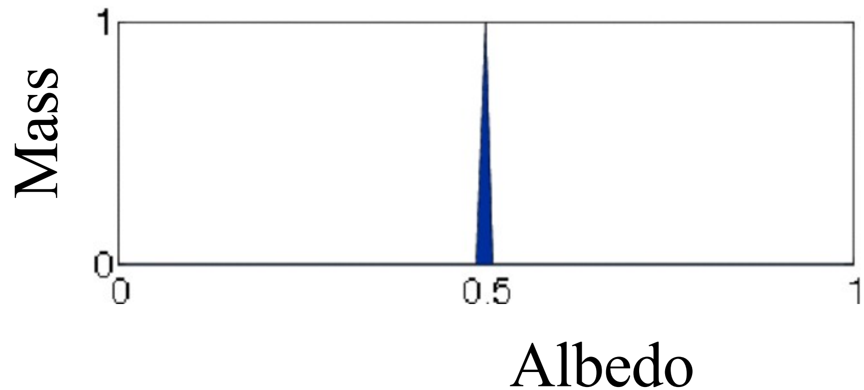

Motivation

- Many objects consist of a small set of dominant colors

- We say such objects satisfy the *k-albedos* constraint

Motivation

- Multi-color objects have have multiple peaks
- A GBR transformation smears the peaks

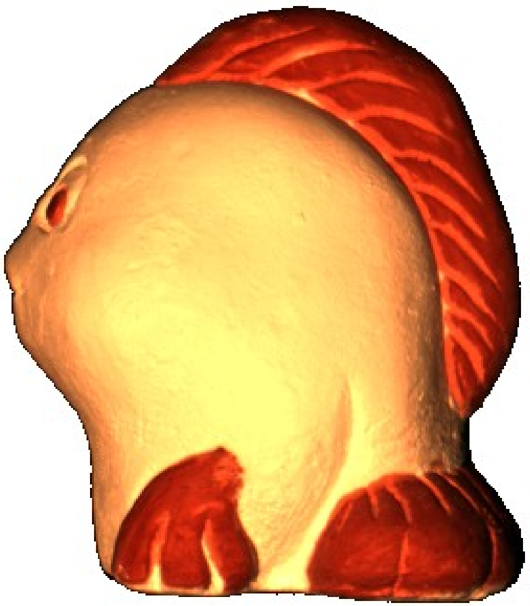

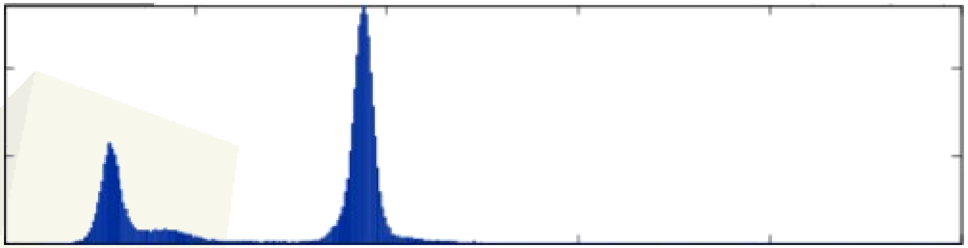

True Albedo Distribution

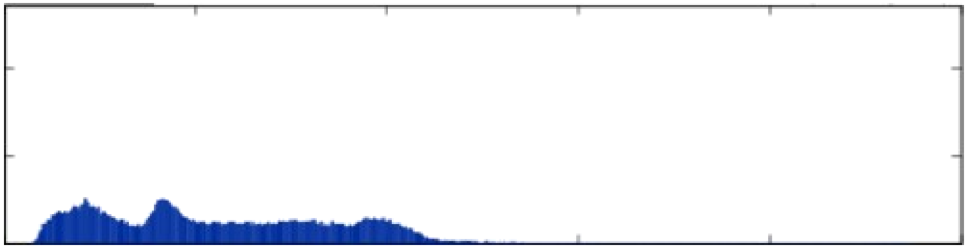

Albedo Distribution Under GBR

Motivation

- Does enforcing "peakiness" resolve the GBR ambiguity?

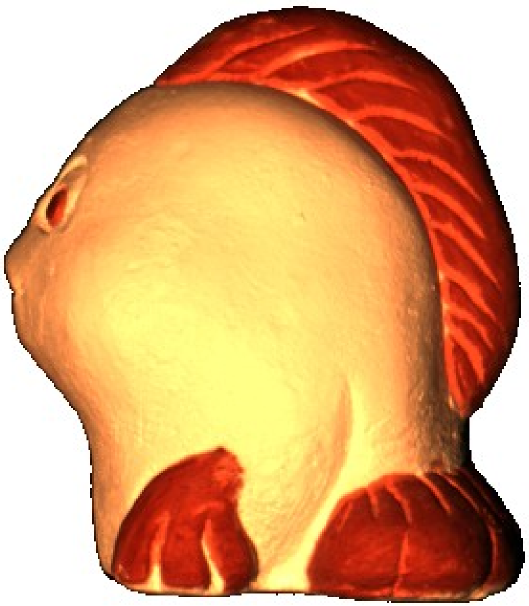

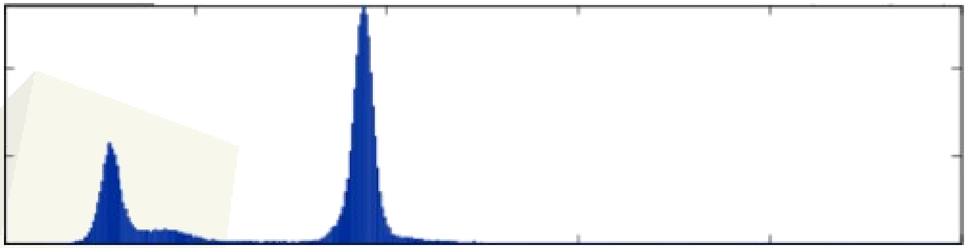

True Albedo Distribution

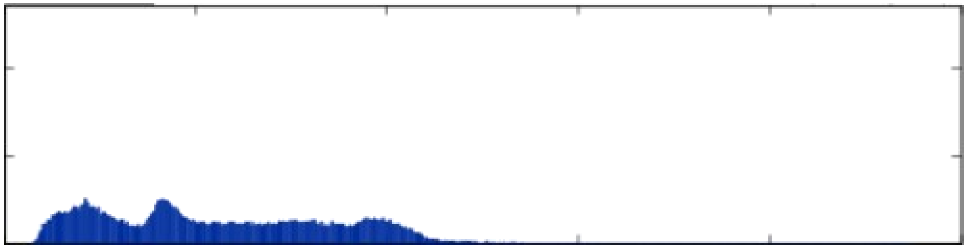

Albedo Distribution Under GBR

Entropy Minimization

- Entropy is a natural measure of the spread of a distribution and conversely its peakiness,

$$H(f) = -\int_S f(x) \log f(x) dx$$

  - f is a p.d.f.
  - S is the support of f

- Doesn't depend on knowledge of *k (# of colors)*
  - Advantage over most clustering algorithms

- Entropy minimization has been successfully applied to other computer vision problems [Finlayson et al. 2004; Palubinskas et al. 1998]

Entropy Minimization

- Let $\tilde{\rho} = \|\hat{\mathbf{b}}^\top \tilde{\mathbf{G}}\| = \|\mathbf{b}^\top \mathbf{G}^{-1} \tilde{\mathbf{G}}\|$ be a random variable representing the albedo distribution obtained by correcting for GBR $\tilde{\mathbf{G}}$.
- Let $f_{\tilde{\rho}}$ be the p.d.f. of $\tilde{\rho}$.
- Objective :
  Minimize $H(f_{\tilde{\rho}}) = -\int_S f_{\tilde{\rho}}(x) \log f_{\tilde{\rho}}(x) dx$
  with respect to $(\tilde{\mu}, \tilde{\nu}, \tilde{\lambda})$ where $\tilde{\mathbf{G}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \tilde{\mu} & \tilde{\nu} & \tilde{\lambda} \end{pmatrix}$

UCSD Jacobs | Computer Science and Engineering

Degenerate Configurations

- *k* albedos on *k* planar surface patches

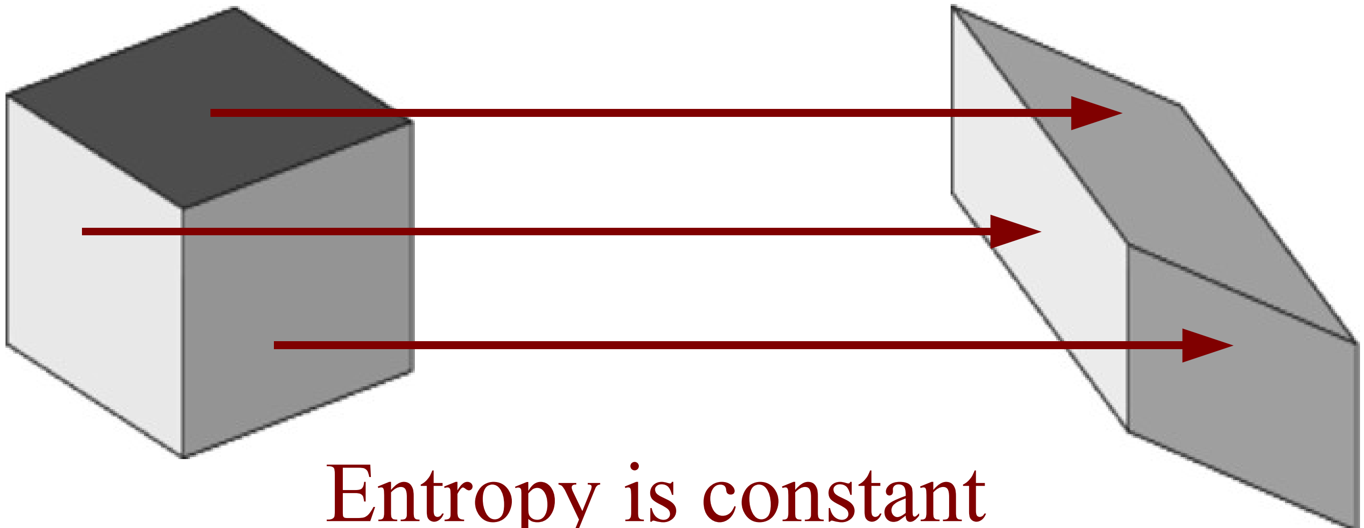

Entropy is constant

$H = \log(3)$

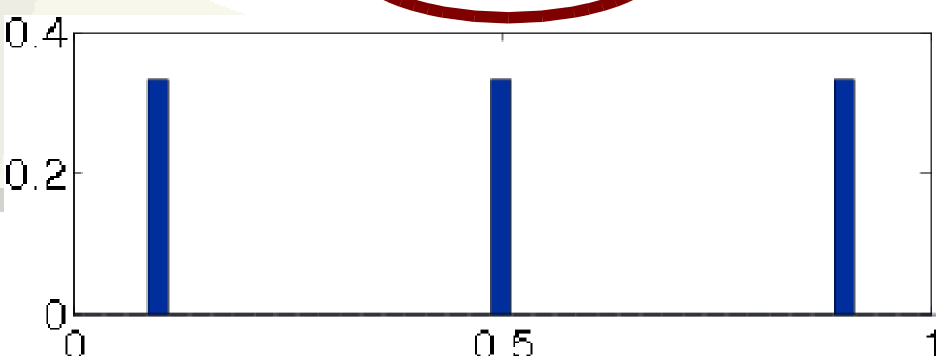

True distribution

$H = \log(3)$

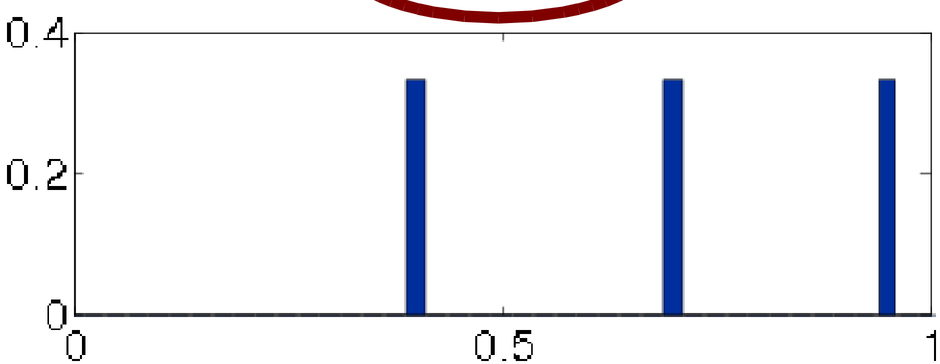

GBR transformed distribution

Degenerate Configurations

- *Nearby albedo modes.*

Entropy is lower under GBR

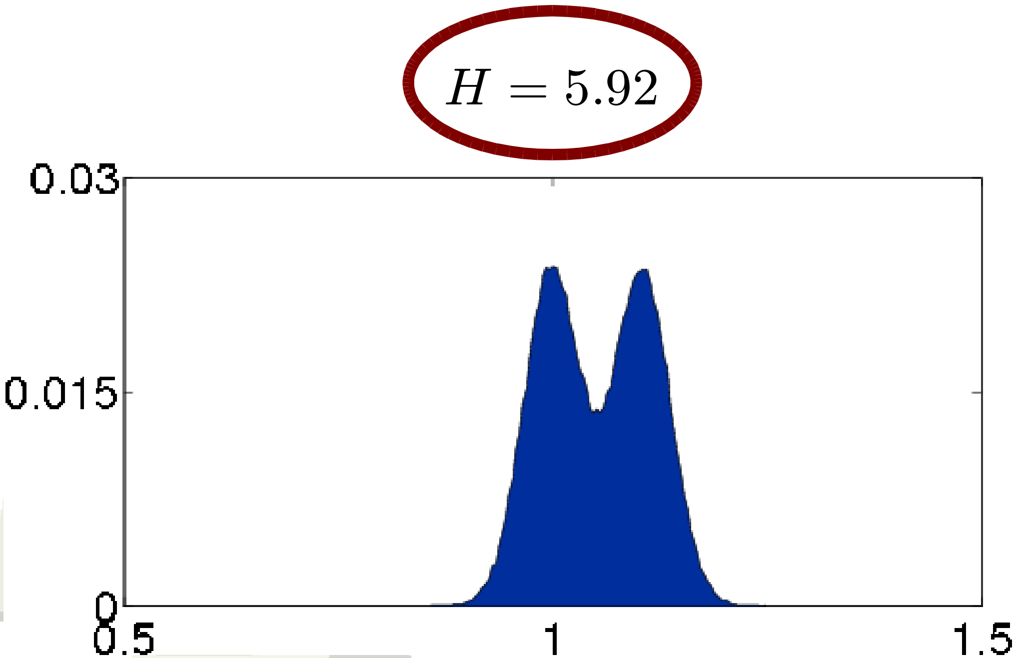

True distribution

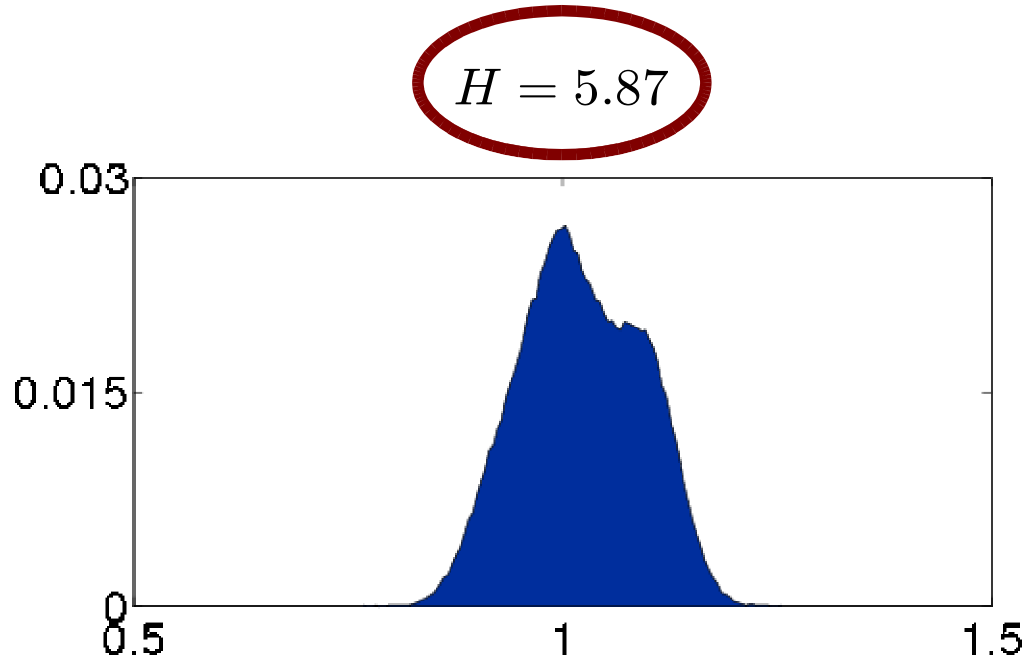

GBR transformed distribution

Theoretical Correctness

- Intuition
  - Each albedo corresponds to a delta function in the distribution
  - A GBR perturbs each mode of the distribution
  - Perturbation has finite support proportional to (μ, ν, λ)
  - Entropy can only increase as long as modes don't overlap

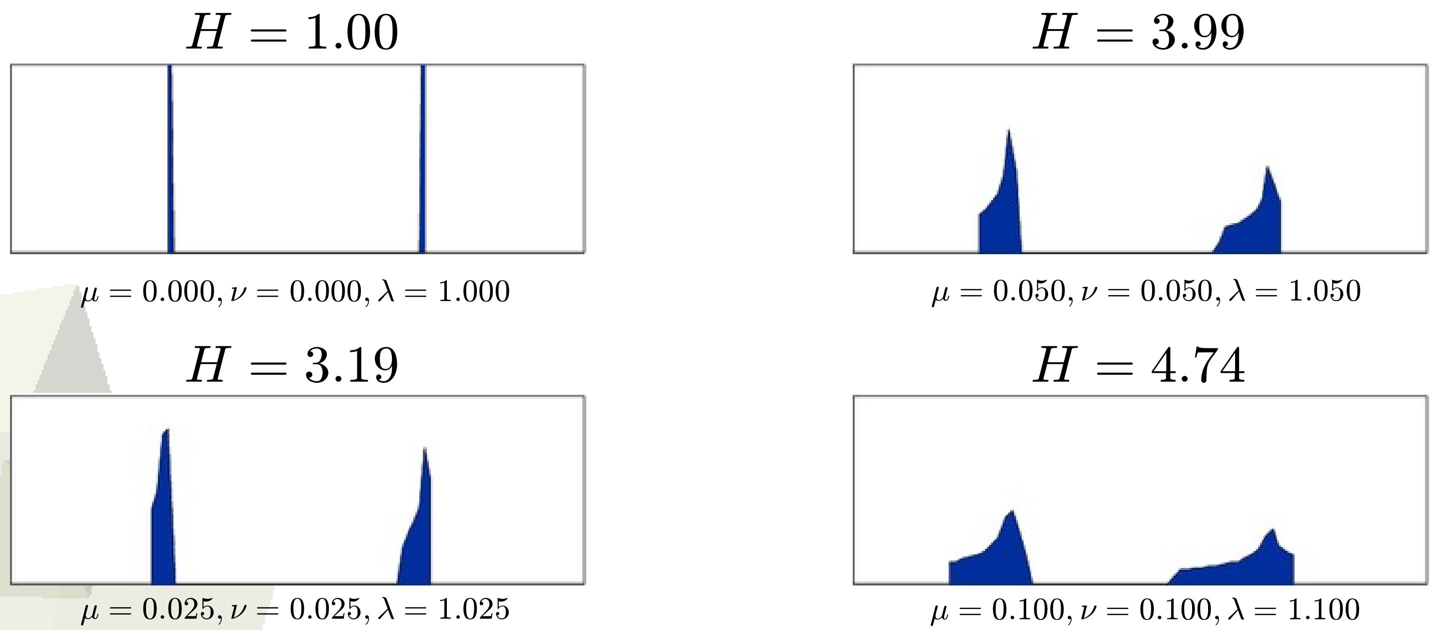

Theoretical Correctness

- Theorem :
  The true GBR parameters (μ, ν, λ) *correspond to a local minima of our objective when the following hold,*
  - *Surface contains* $k \ll N$ *albedo values*
  - *Non-degenerate surface normal distribution*
  - *Albedo is independent of surface normal*

- Sufficient (not necessary) conditions.
- More details in paper.

Approximating the Entropy

- Entropy is defined w.r.t. a density function
- In practice only have samples drawn from the distribution
- Options :
  - Fit a continuous distribution (using Parzen windows, mean-shift algorithm, kernel based estimators, etc.) then compute entropy
  - Approximate entropy from a histogram

Approximating the Entropy

- MLE estimator of entropy,

$$\hat{H}_{MLE}(f) = -\sum_{i=1}^{m} \frac{a_i}{n} \log \frac{a_i}{n}$$

  - $\{a_i\}, i = 1...m$ is an m-bin histogram computed from n samples
- $\hat{H}_{MLE}$ is biased, but variance is low
- $m = 256$ used in our experiments

Algorithm Overview

- Input : M images of N pixels each
- Algorithm
  i. Recover normals / albedos up to a GBR using algorithm of Yuille and Snow
  ii. Perform discrete search over GBR parameters μ, ν, λ
    a) Apply GBR
    b) Compute m-bin histogram $\{a_i\}, i = 1...m$ of albedo values
    c) Approximate entropy using

$$H(f_\rho) \simeq -\sum_{i=1}^{m} \frac{a_i}{N} log \frac{a_i}{N}$$

  iii. Apply GBR corresponding to $\arg\min_{\mu,\nu,\lambda} H(f_\rho)$
  iv. Integrate to obtain surface

UCSD Jacobs | Computer Science and Engineering

Experimental Results

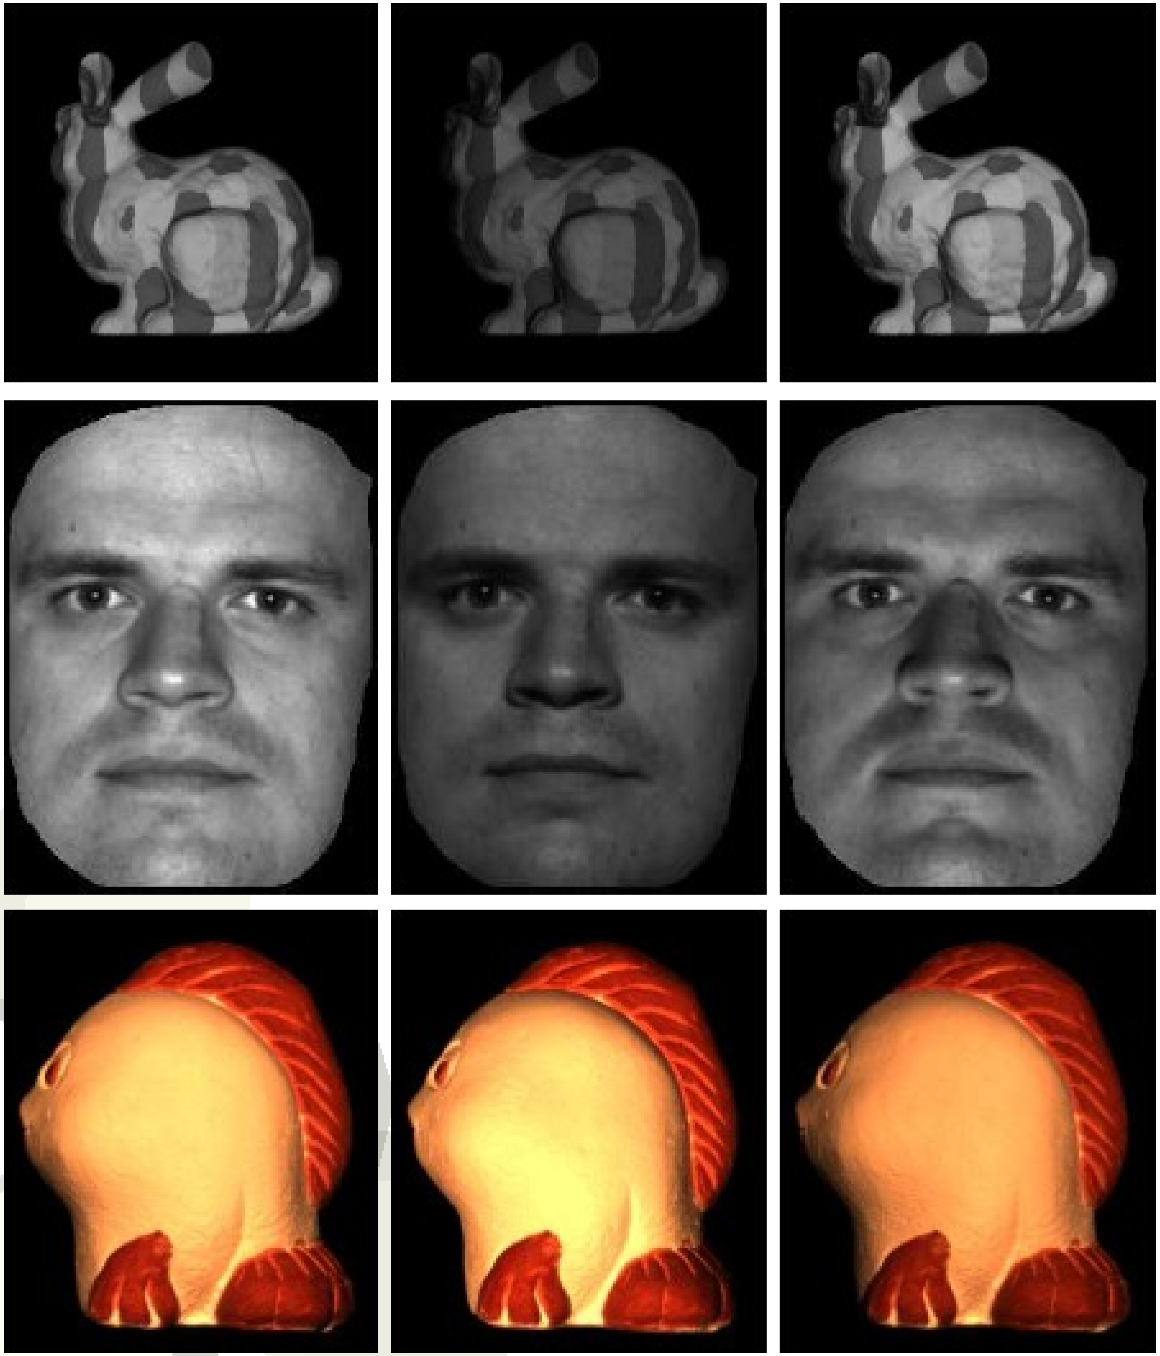

- Input Images
- Three datasets

Stanford Bunny (Synthetic)

6 input images

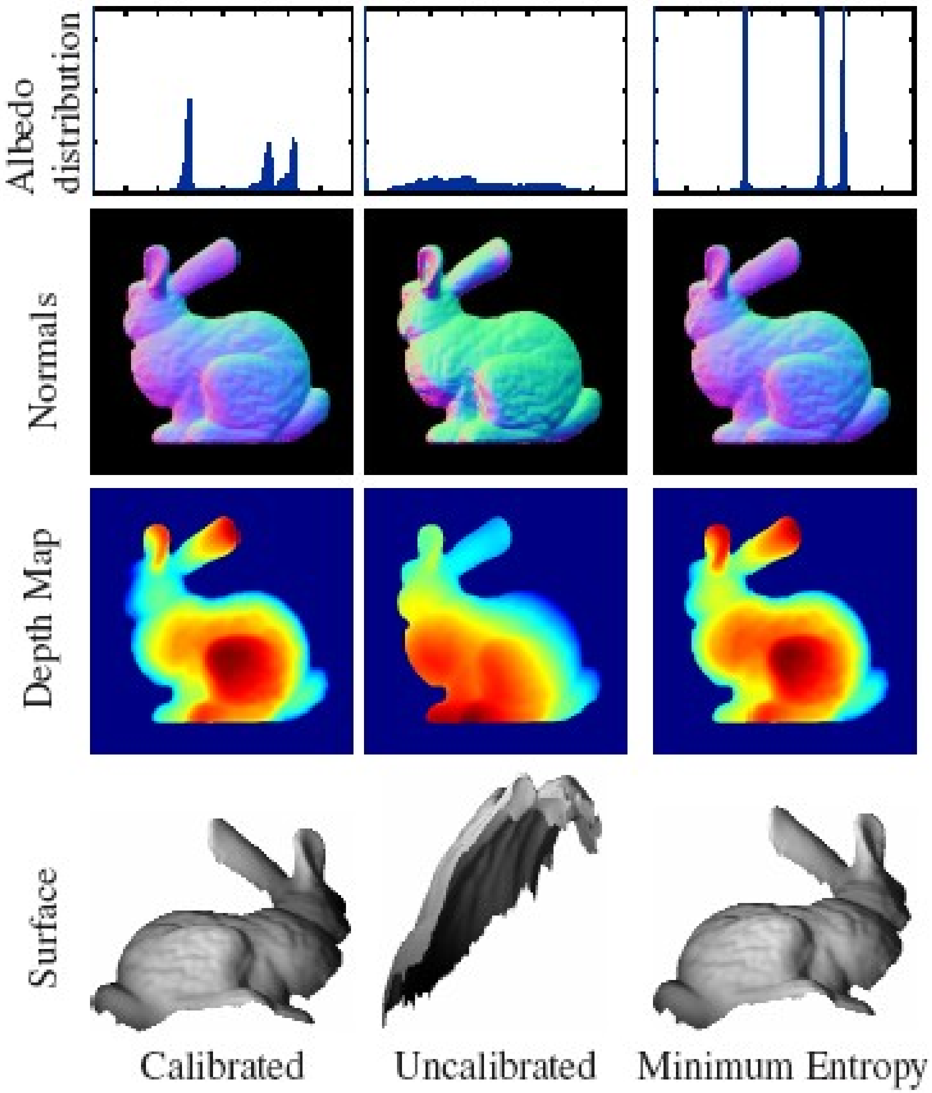

Yale B Face Database

9 input images

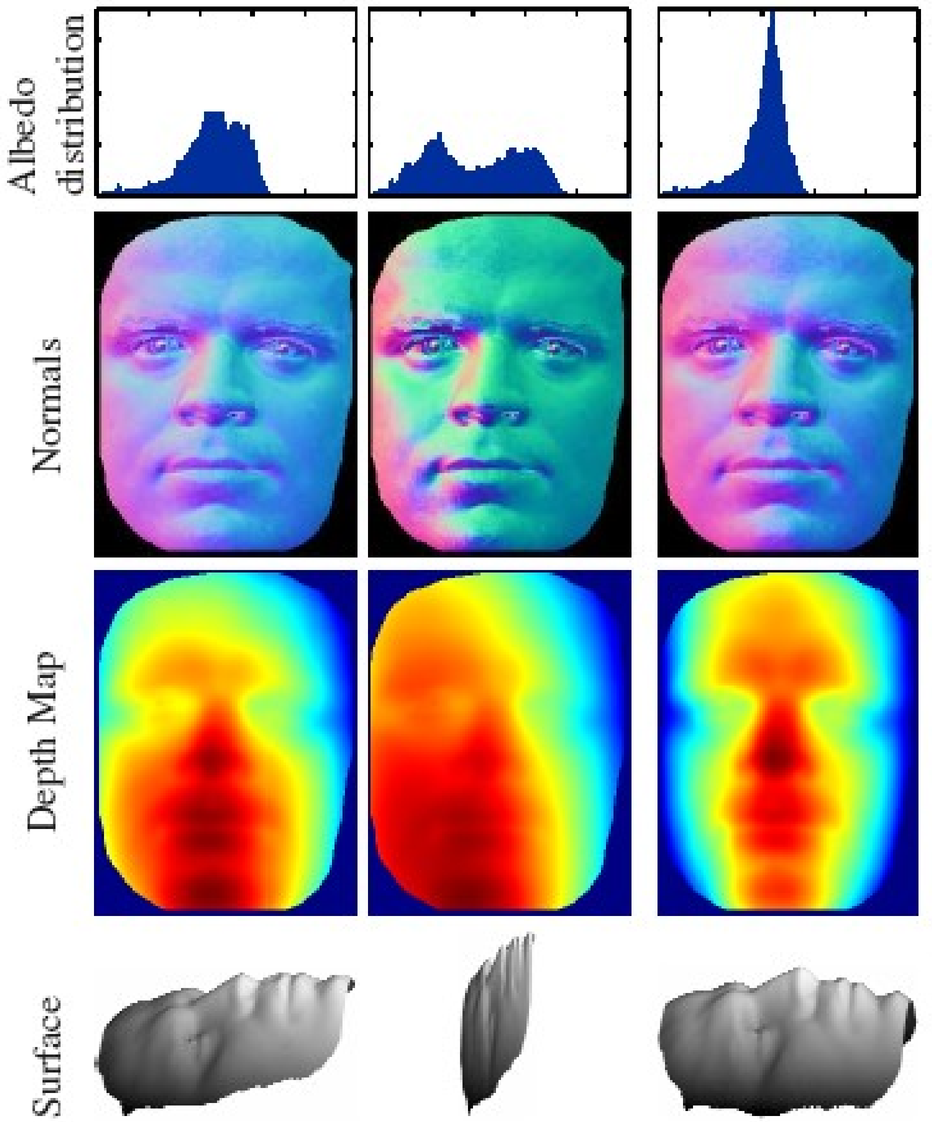

Red & Yellow Fish

5 input images

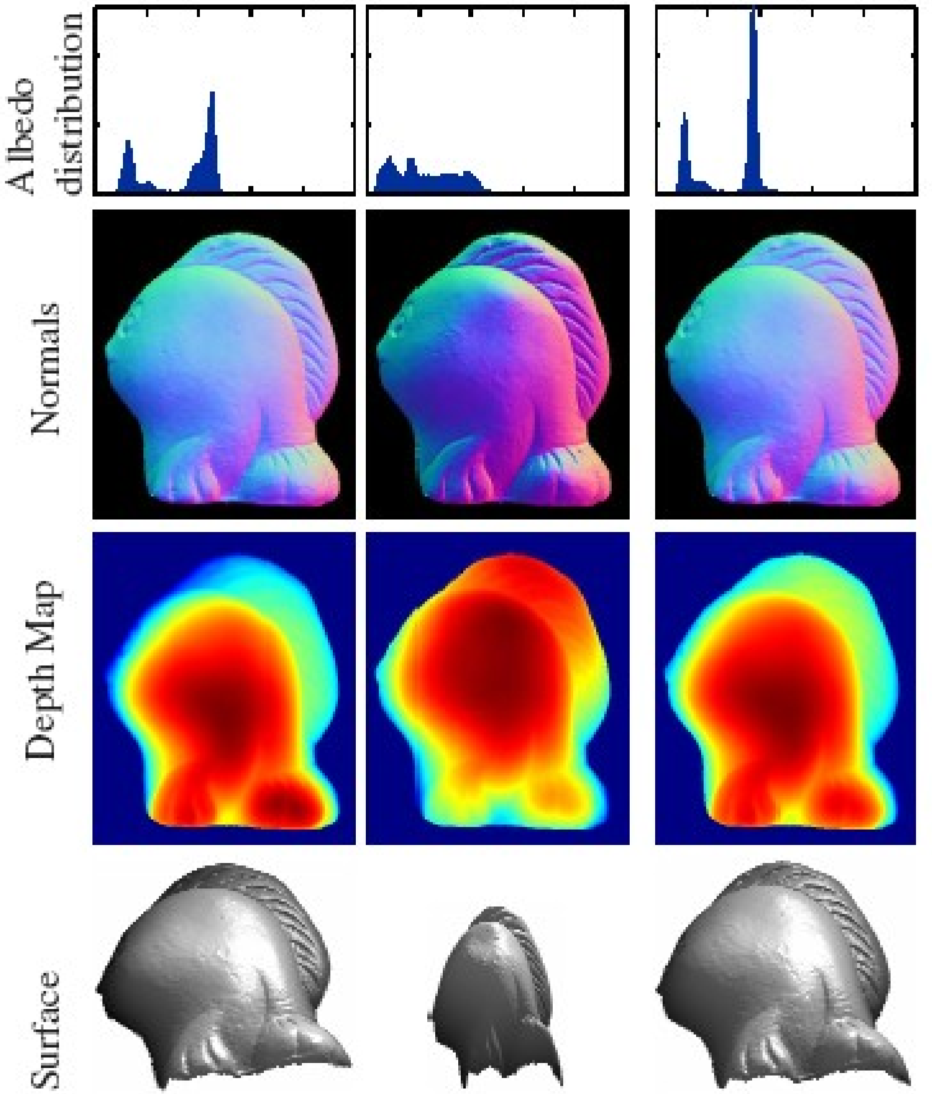

UCSD Jacobs | Computer Science and Engineering

Review

- Novel constraint
  - Albedo distribution should have low entropy
  - Valid for large class of real-world objects

- New method to resolve the GBR ambiguity
  - By minimizing entropy of albedo distribution

- Validation
  - Theoretical
  - Experimental

Questions?

9 input images

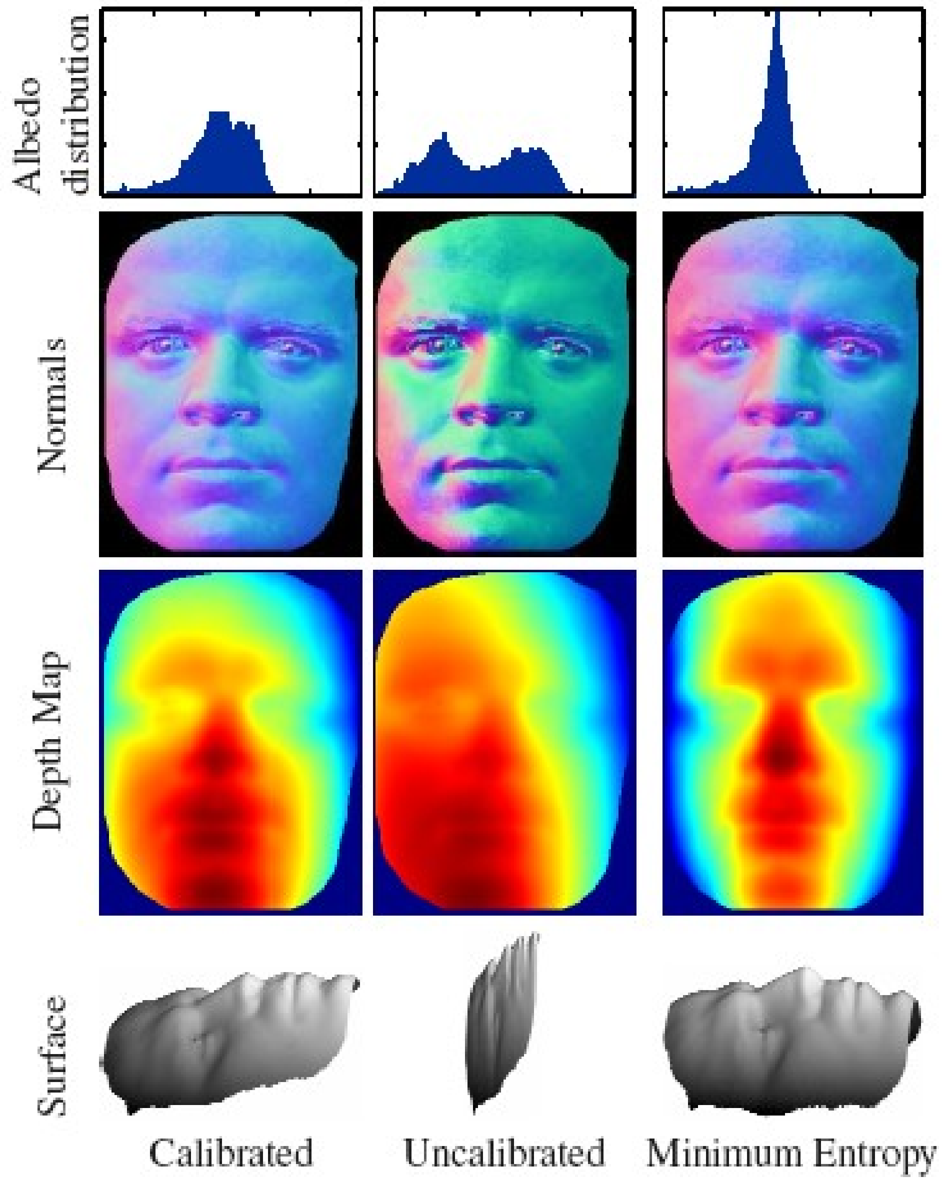

UCSD Jacobs | Computer Science and Engineering

Future Work

- Exploit color information
  - Increases separability in the albedo distribution
- Exploit geometry
  - Facet vectors ($\hat{\mathbf{b}} = \hat{\rho}\hat{\mathbf{n}}$) should lie on concentric ellipsoids under a GBR

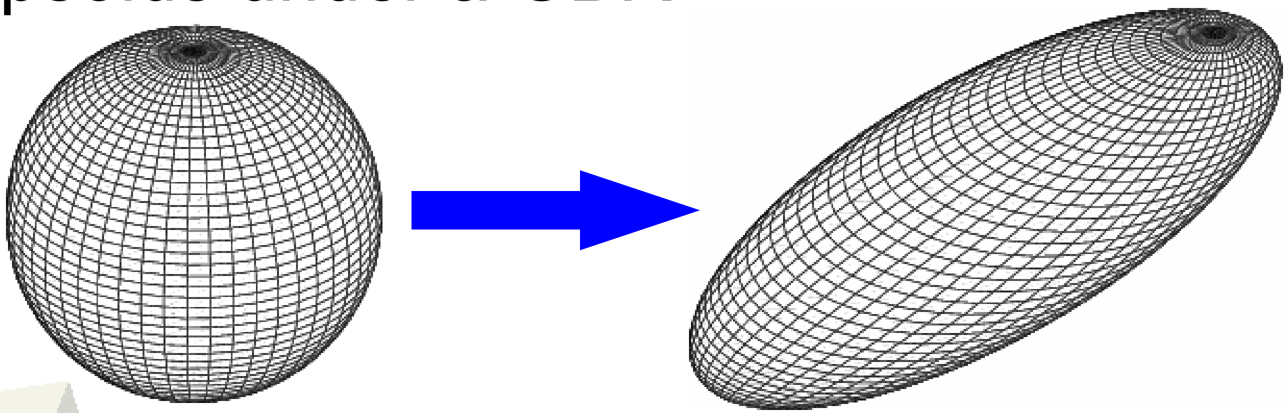

- Apply method to other ambiguities
  - Lorentz ambiguity [Basri & Jacobs 2001]
  - KGBR ambiguity [Yuille et al. 2001]