Computational Photography and Video: Intrinsic Images

Prof. Marc Pollefeys
Dr. Gabriel Brostow

ETH
inf | Informatik Computer Science

Last Week

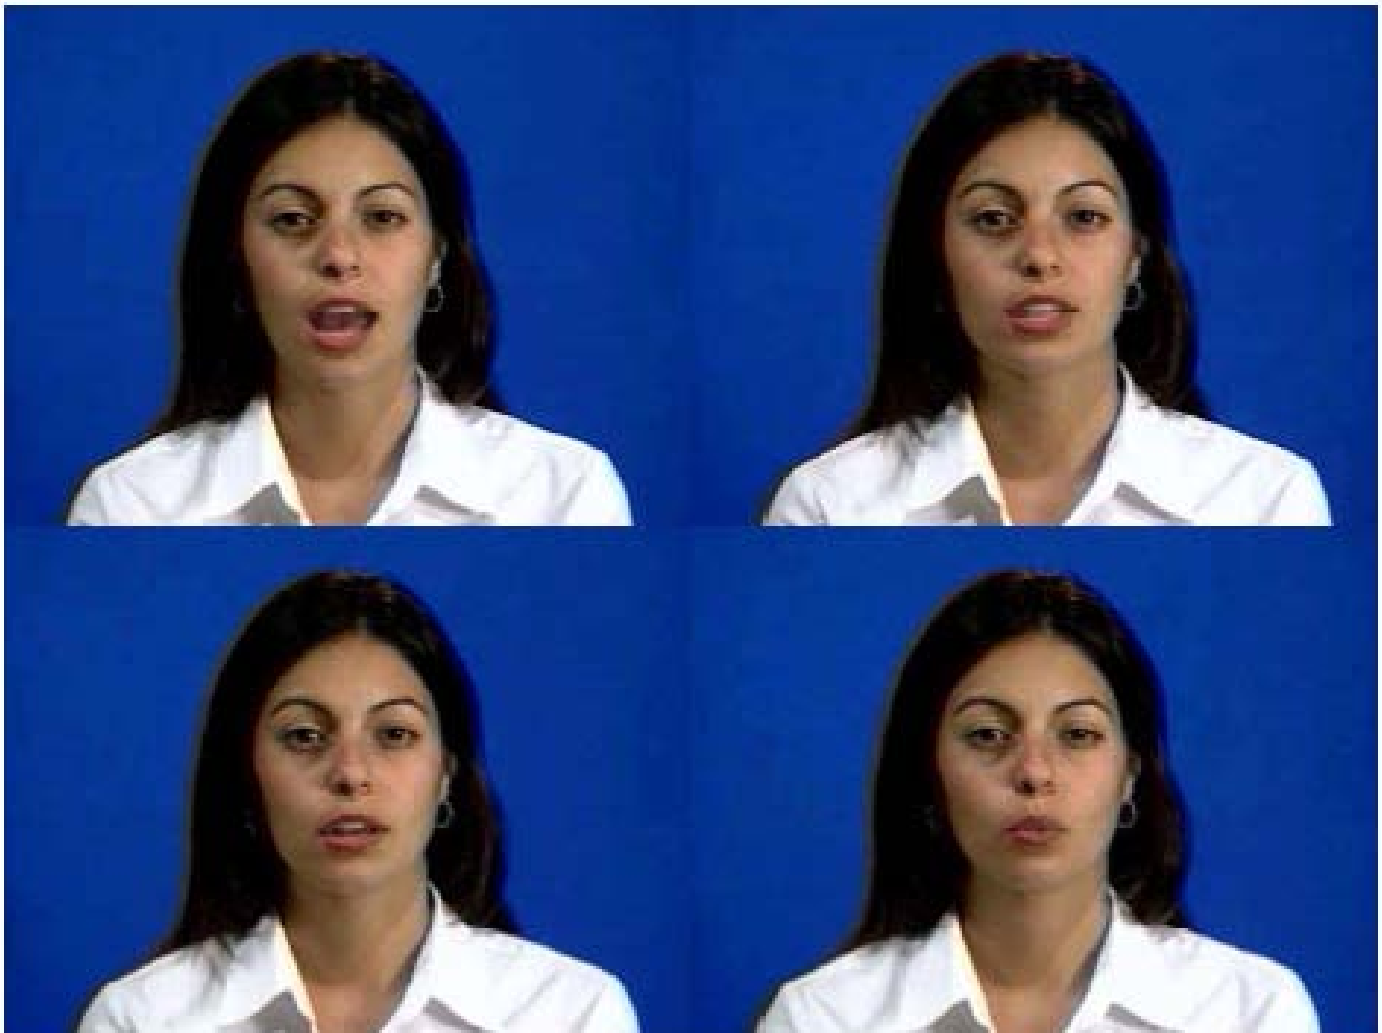

Schedule	Computational Photography and Video	Exercises
18 Feb	Introduction to Computational Photography	
25 Feb	More on Cameras, Sensors and Color	Assignment 1: Color
4 Mar	Warping, morphing and mosaics	Assignment 2: Alignment
11 Mar	Image pyramids, Graphcuts	Assignment 3: Blending
18 Mar	Dynamic Range, HDR imaging, tone mapping	Assignment 4: HDR
25 Mar	Video Synthesis I	Papers
1 Apr	Video Synthesis II	Papers
8 Apr	Intrinsic Images	Project proposals
15 Apr	*Easter holiday – no classes*	
22 Apr	Vectorizing Rasters	Papers
29 Apr	Non-photorealistic Rendering & Animation	Project updates
6 May	Time-Lapse Video	Papers
13 May	(Re)Coloring	Papers
20 May	Video Based Rendering	Optional team meetings
27 May	Final Project Presentations I	Project Presentations II

ETH

Example Problem: Background Normalization

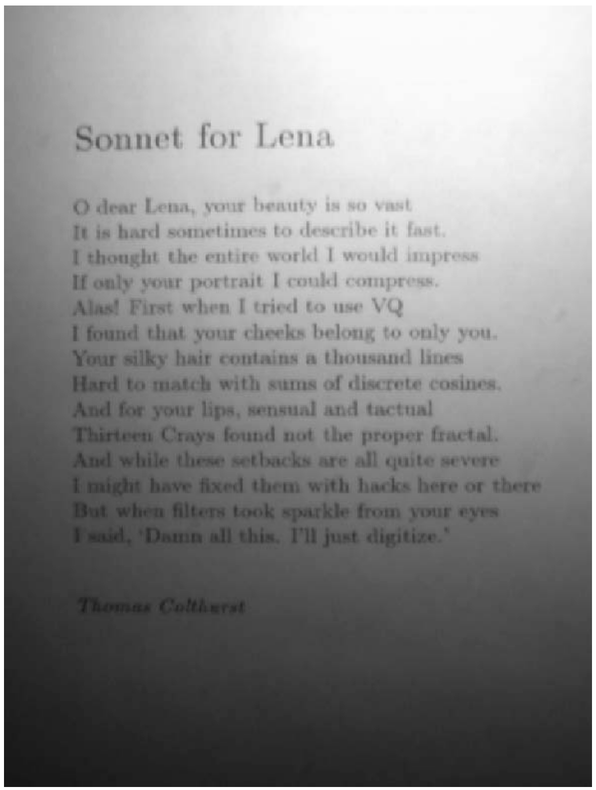

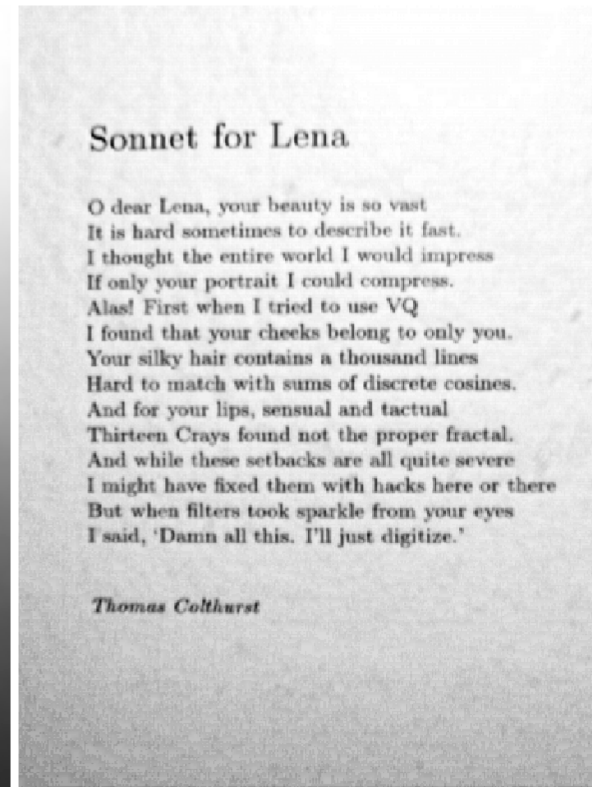

When taking a picture, what color is a (Lambertian) surface?



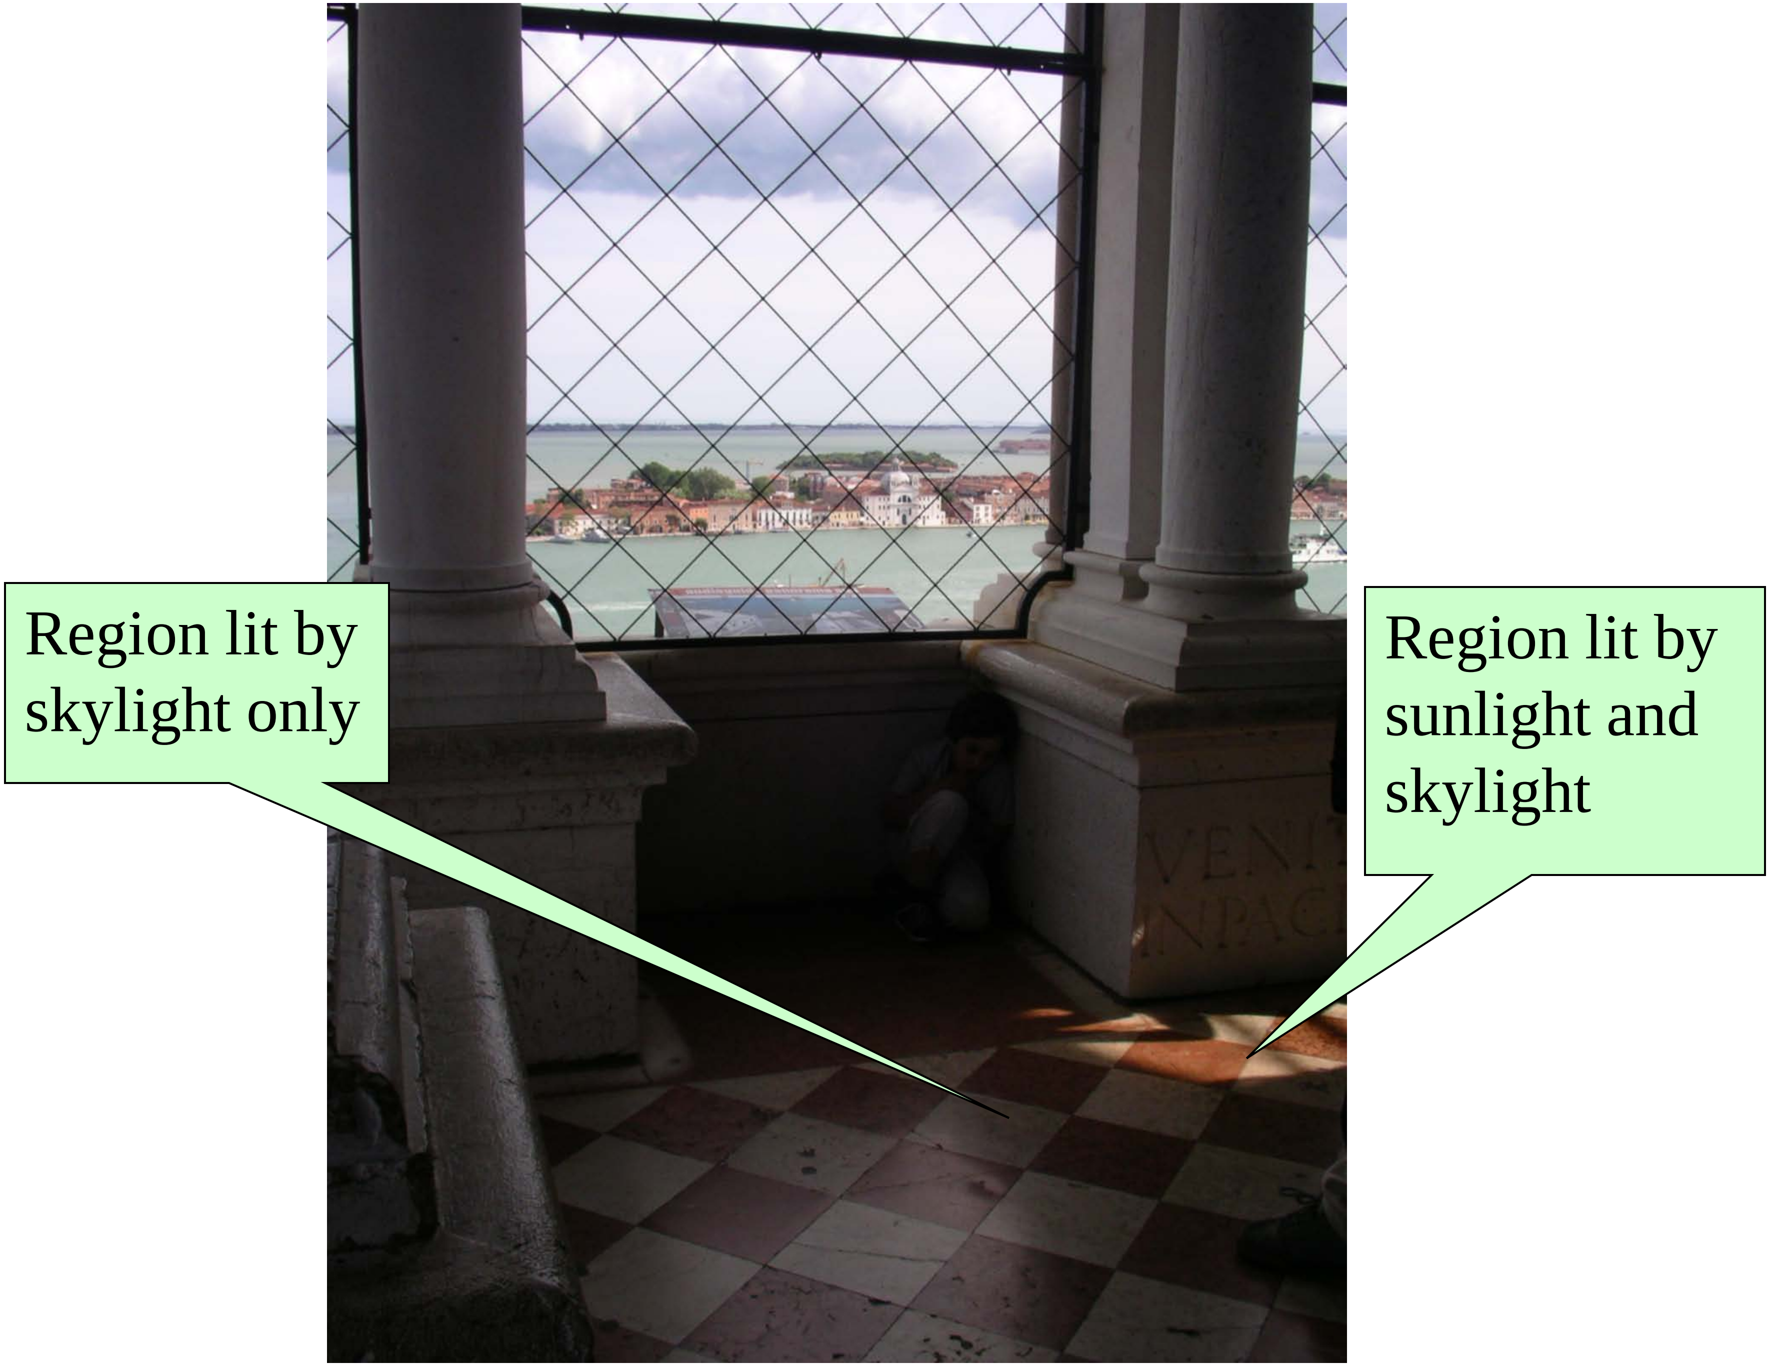
Region lit by skylight only
Region lit by sunlight and skylight

What great things could we do if we could easily find shadows?



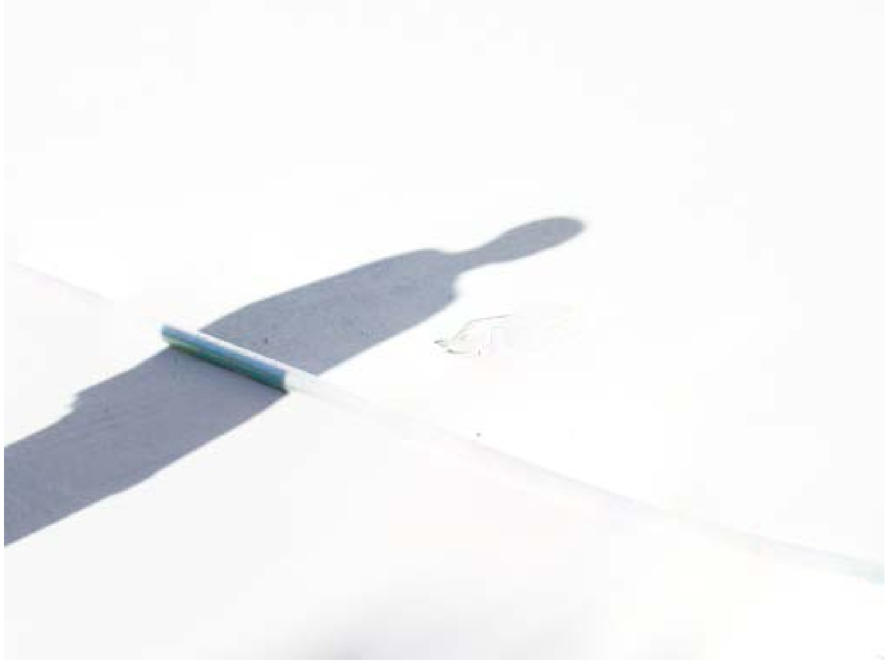

An Intrinsic Image

- What effect is the lighting having, irrespective of surface materials?

- What is the surface reflectance, irrespective of lighting?

Tappen et al. PAMI'05

Original

Lighting/Shading

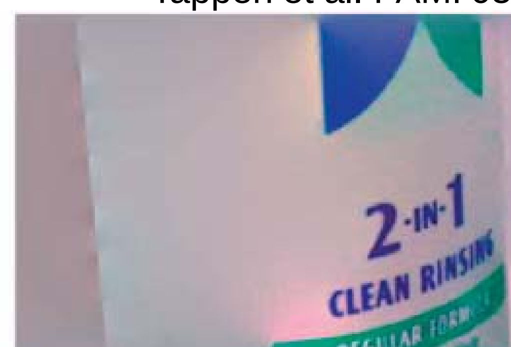

Reflectance

ETH

Pursuit of Intrinsic Images (1)

- Lightness and Retinex Theory
  - Land & McCann ’71

- Recovering Intrinsic Scene Characteristics From Images
  - Barrow & Tenenbaum ‘78

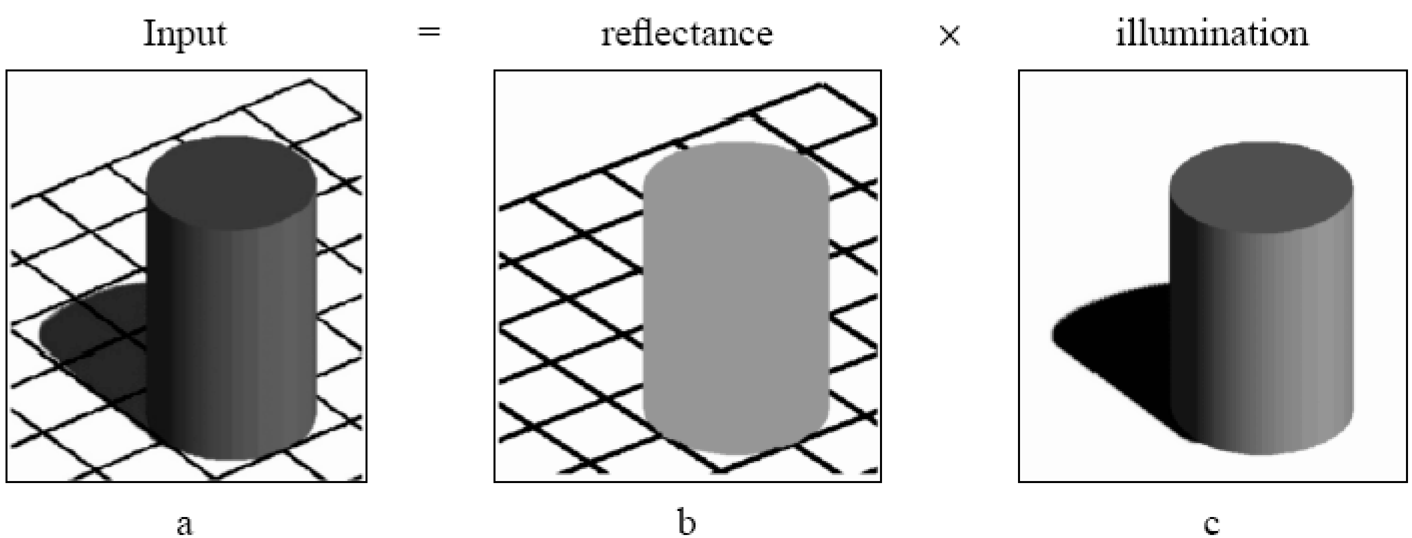

Pursuit of Intrinsic Images (2)

- Painted Polyhedra - ICCV’93
- Image Sequences - ICCV’01
- Single Image - NIPS’03
- Entropy Minimization - ECCV’04

Pursuit of Intrinsic Images (2)

- Painted Polyhedra - ICCV’93 (Generative)
- Image Sequences - ICCV’01 (Discriminative)
- Single Image - NIPS’03 (Discriminative)
- Entropy Minimization - ECCV’04 (Generative)

I

Painted Polyhedra

- Recovering Reflectance and Illumination in a World of Painted Polyhedra
  - Sinha & Adelson, ICCV’93

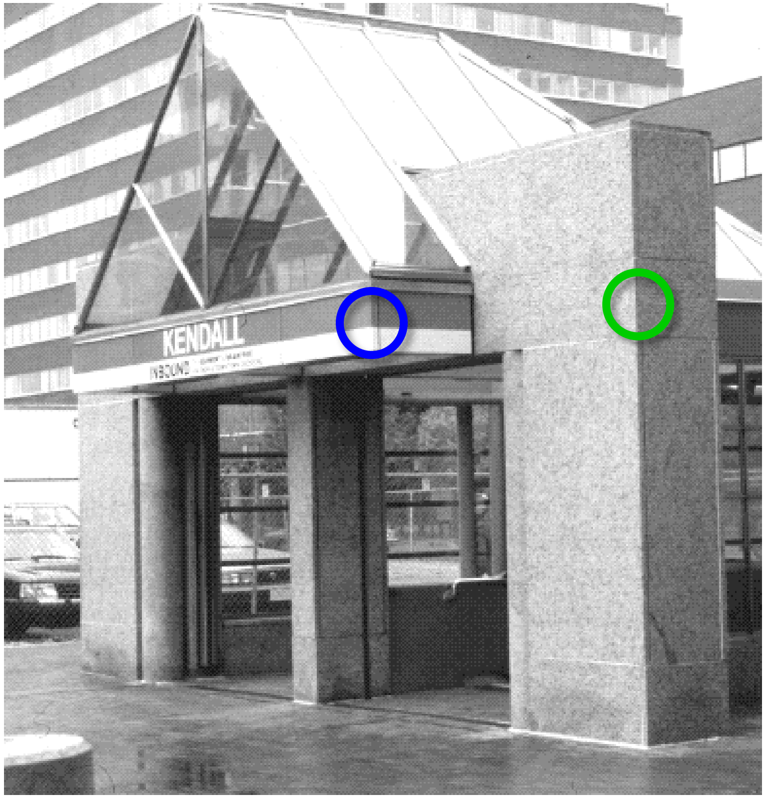

ETH

Not All Edges are Equal

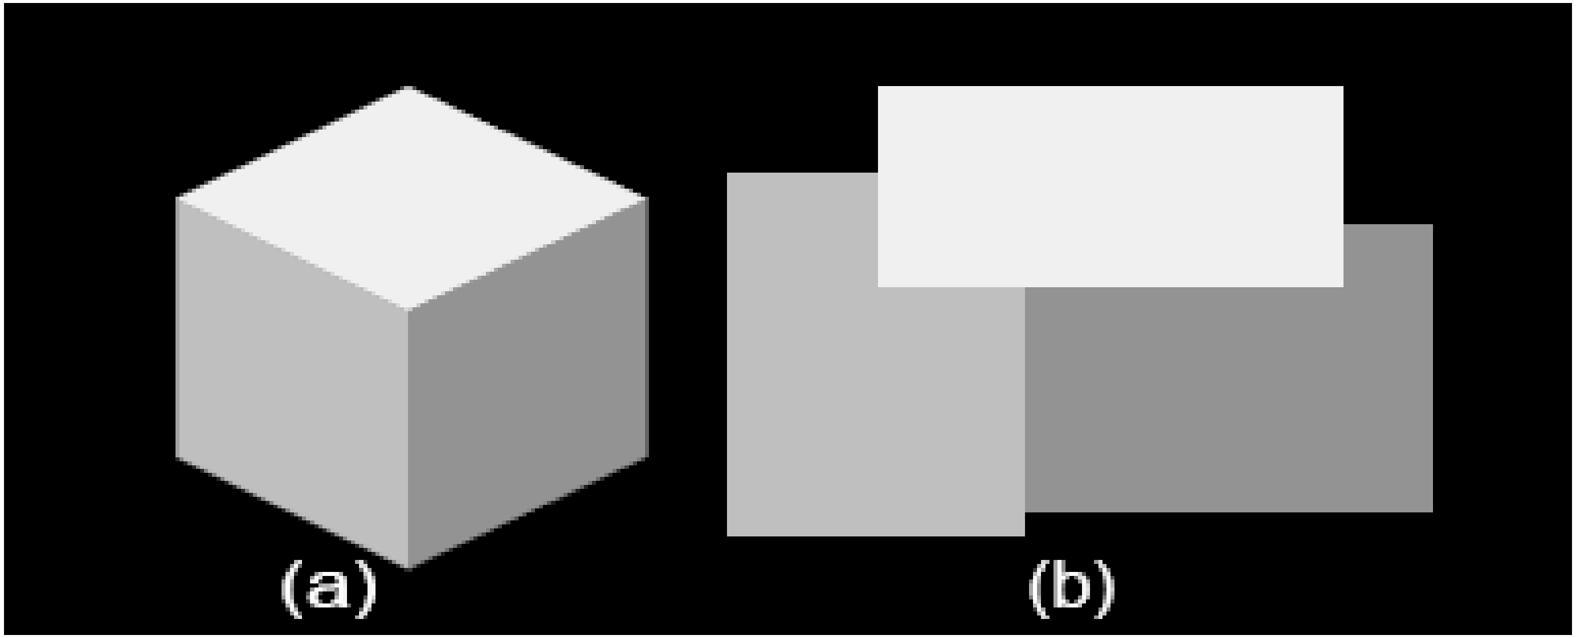

Local Edges are a Hint?

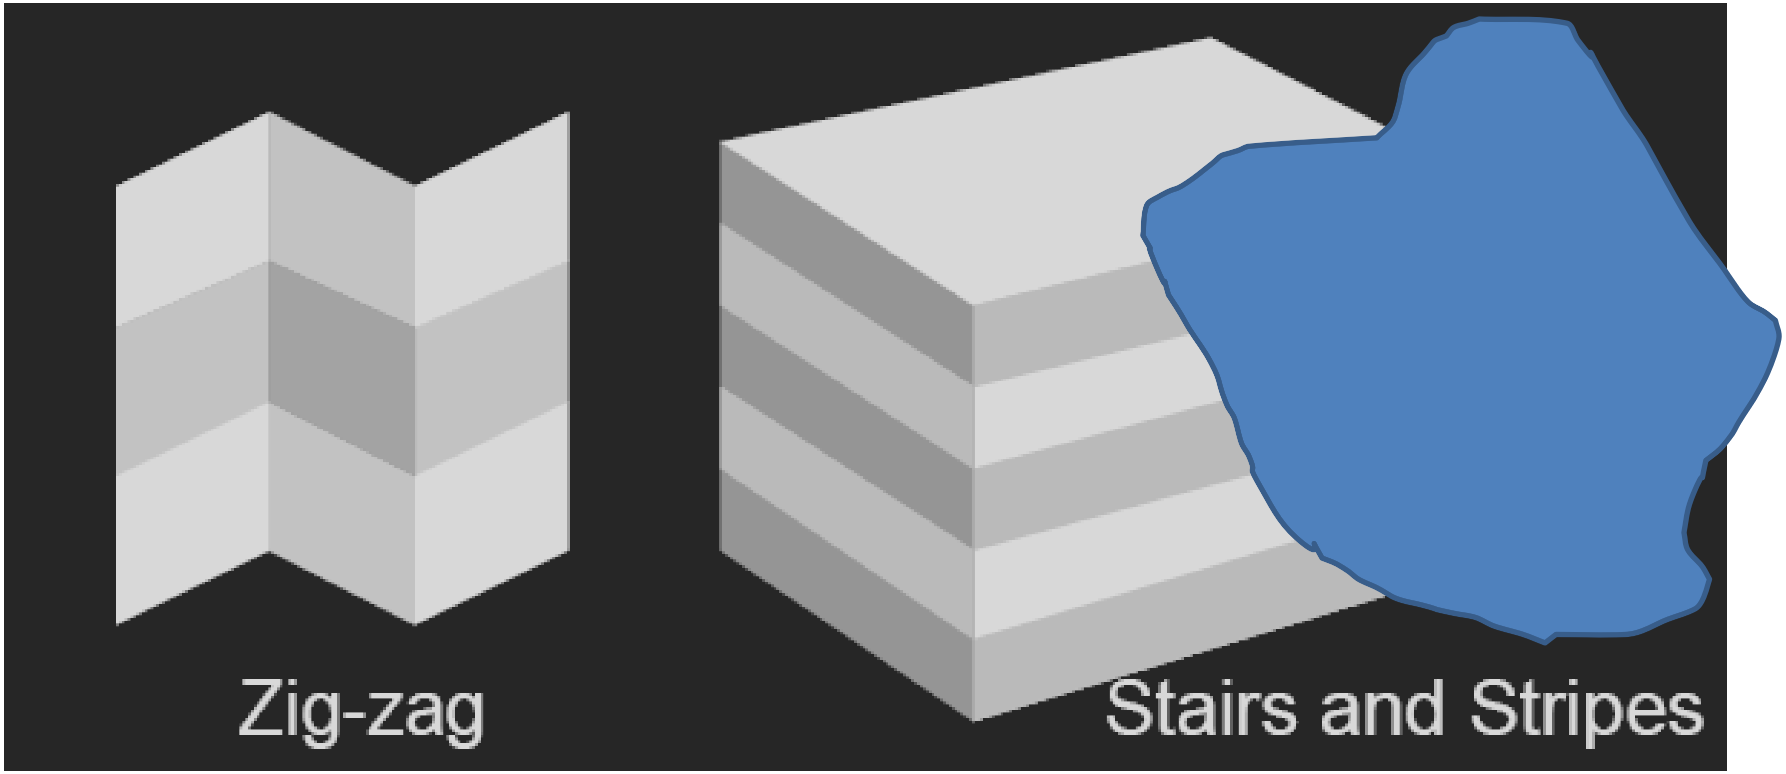

Edge Junctions are Useful

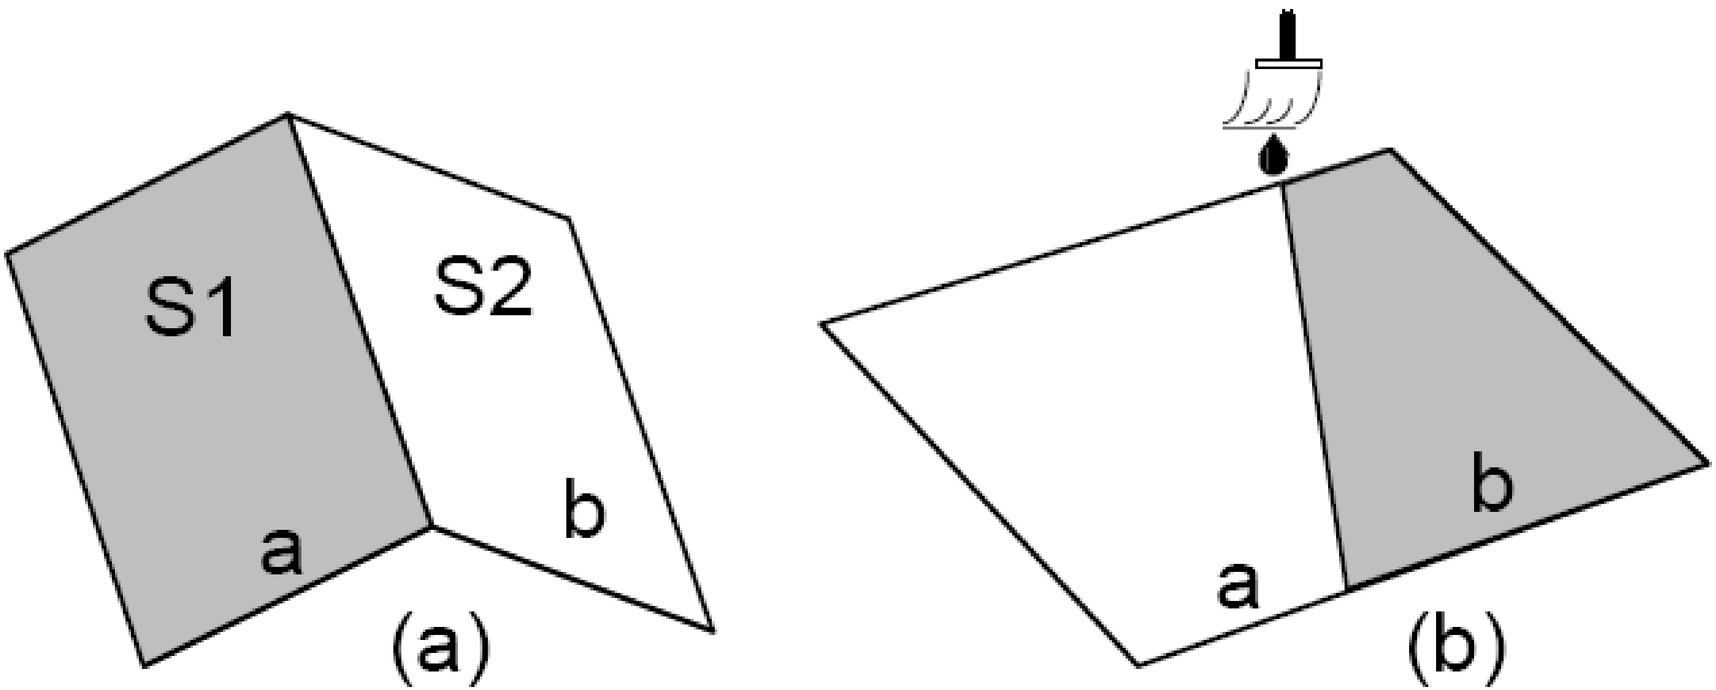

Junction Catalog

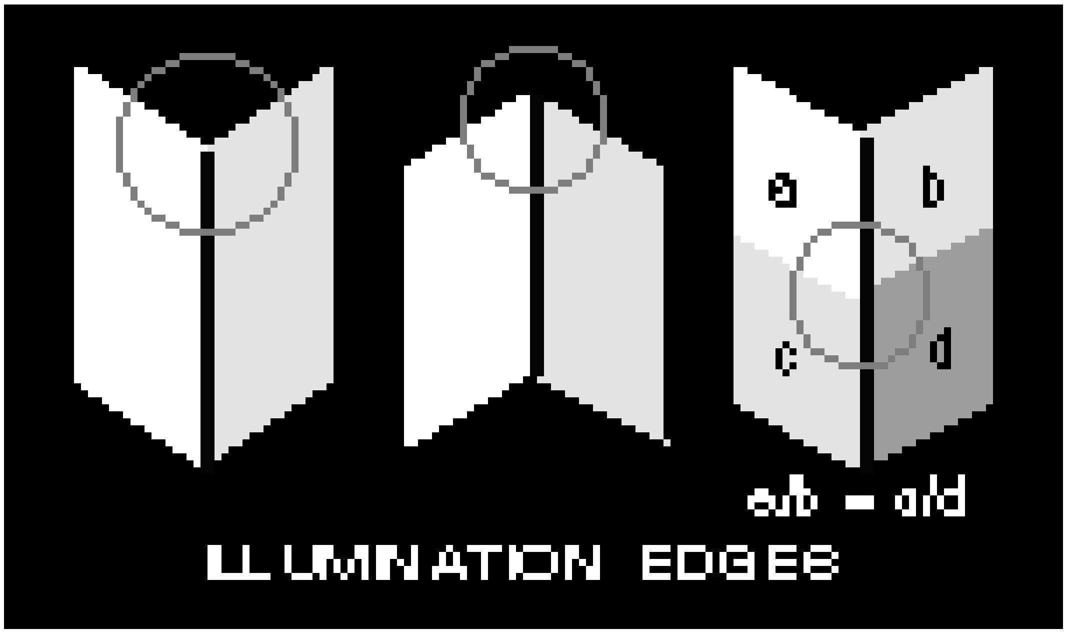

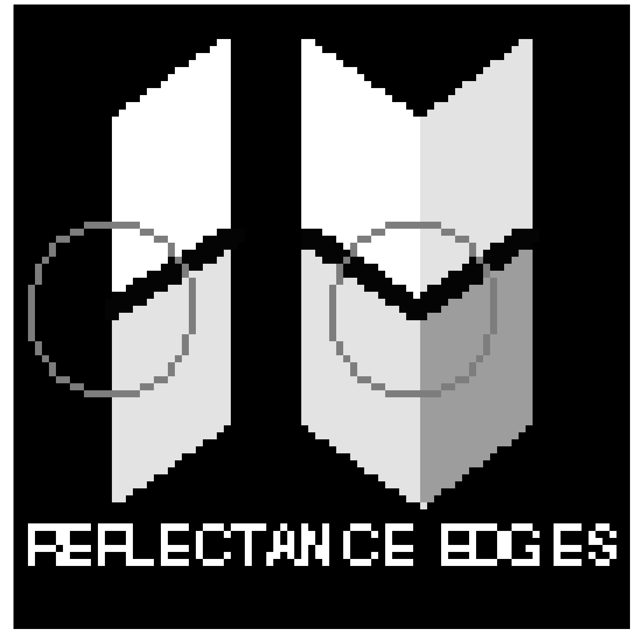

‘Y’, ‘arrow,’ and ‘psi’ junctions

‘T’ junctions

ETH

Examples

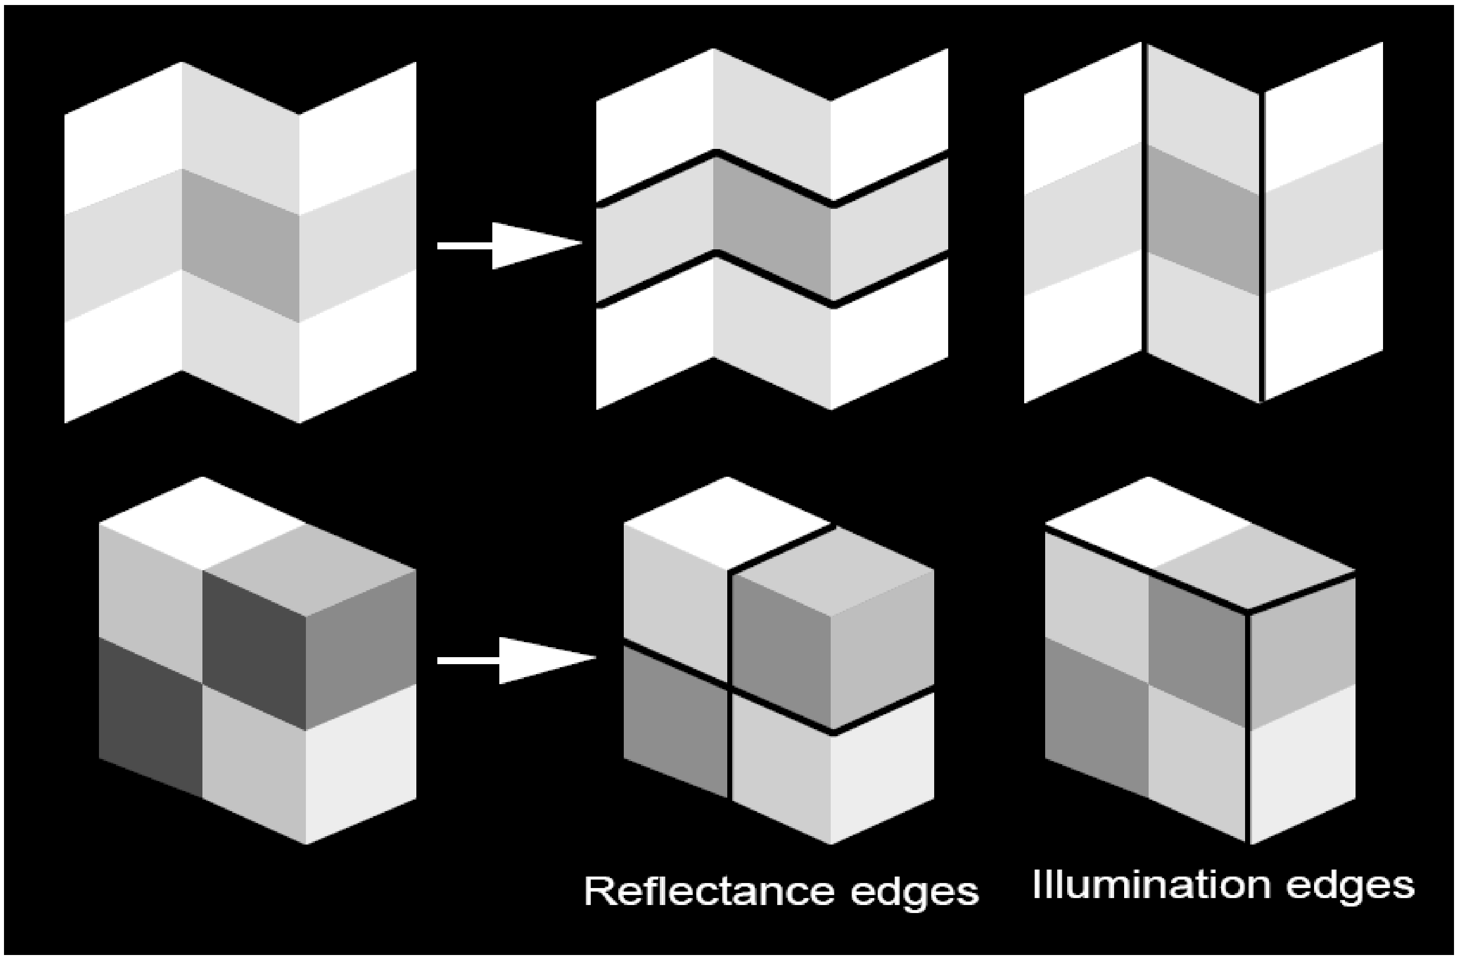

Examples

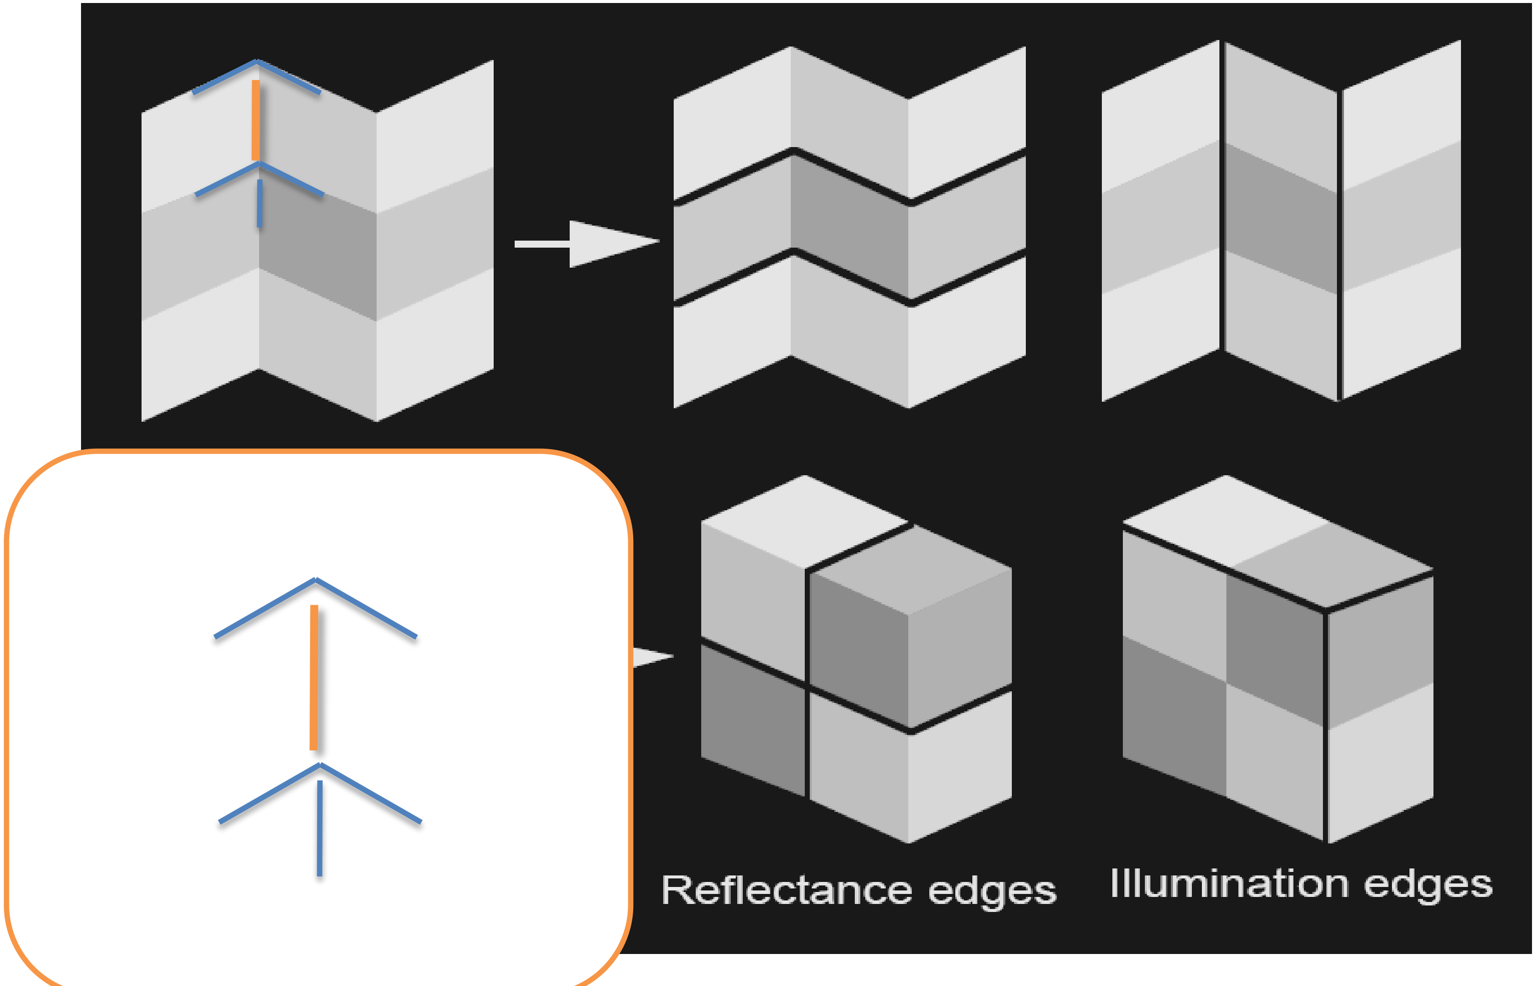

Examples

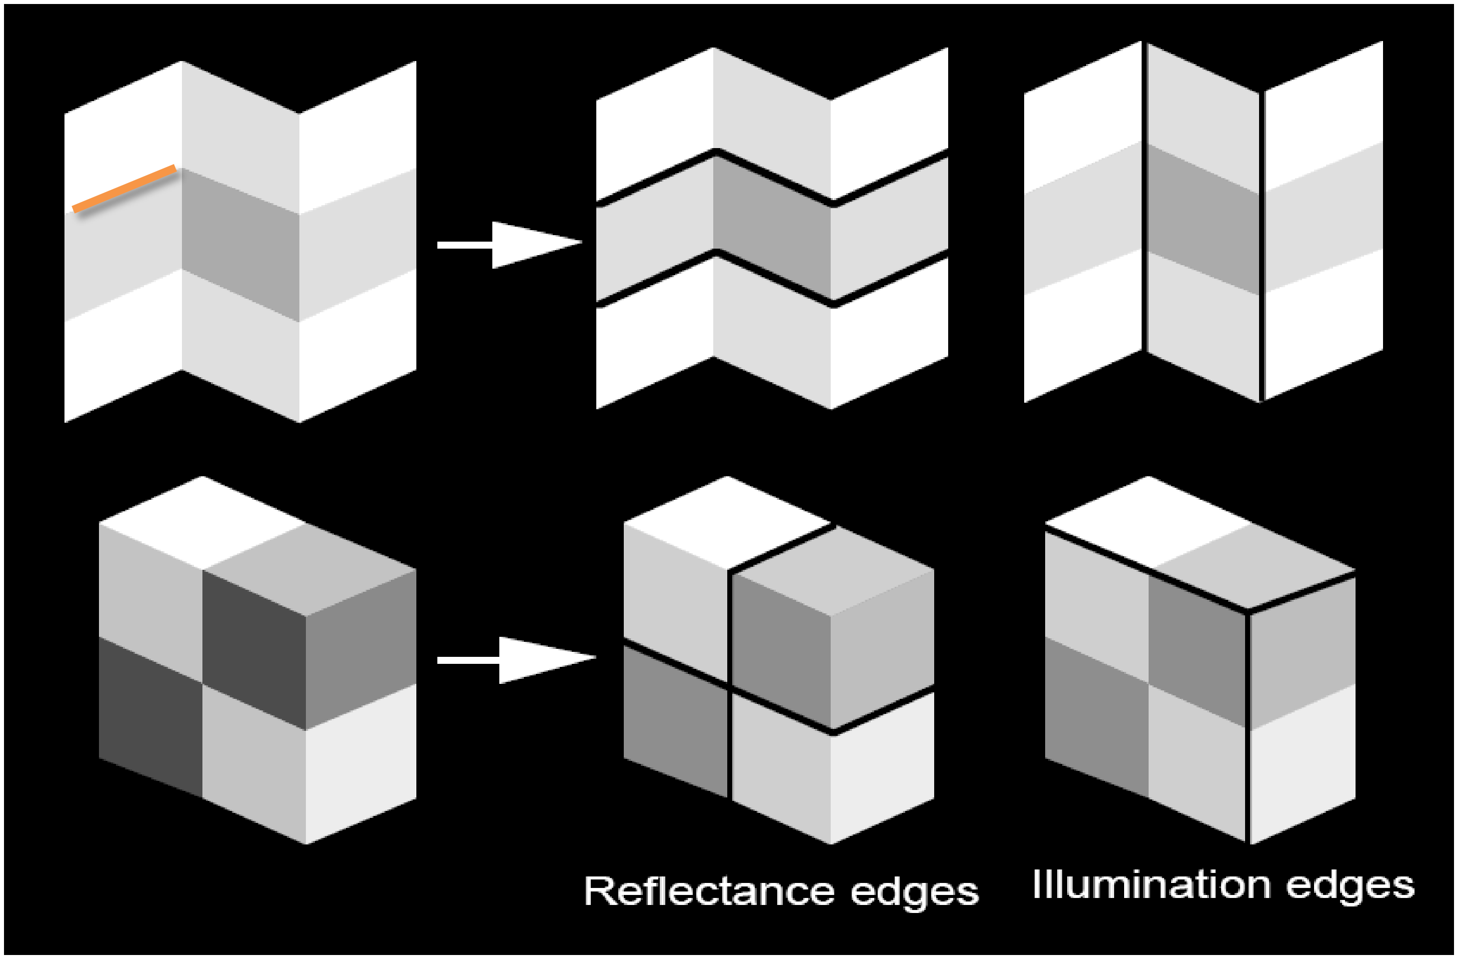

Examples

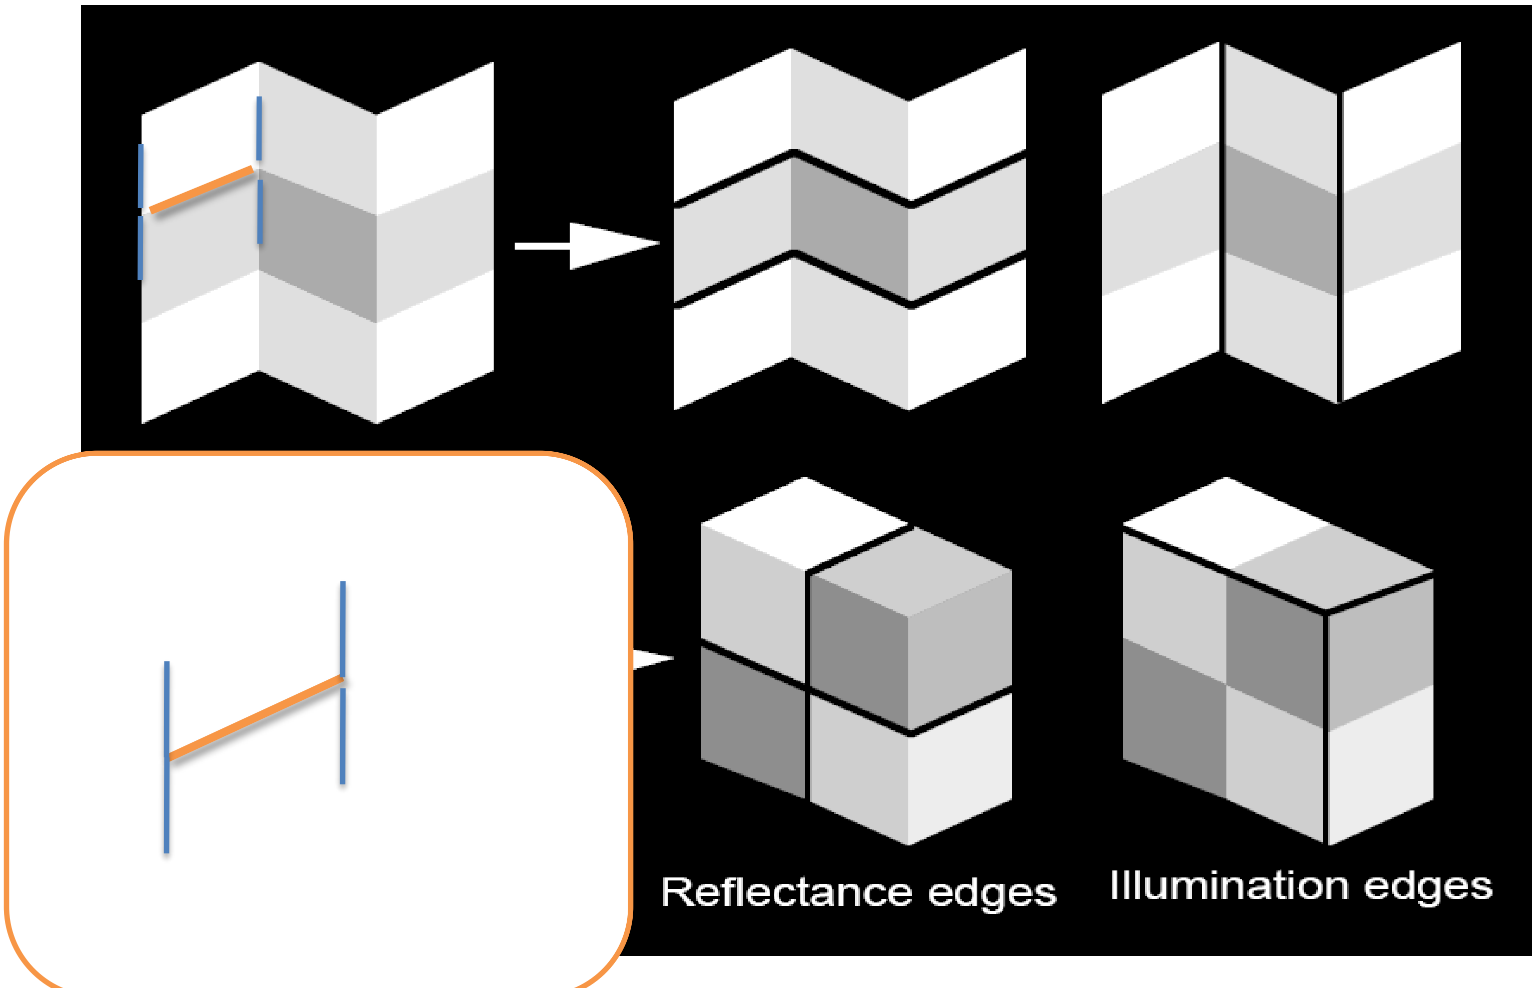

Examples

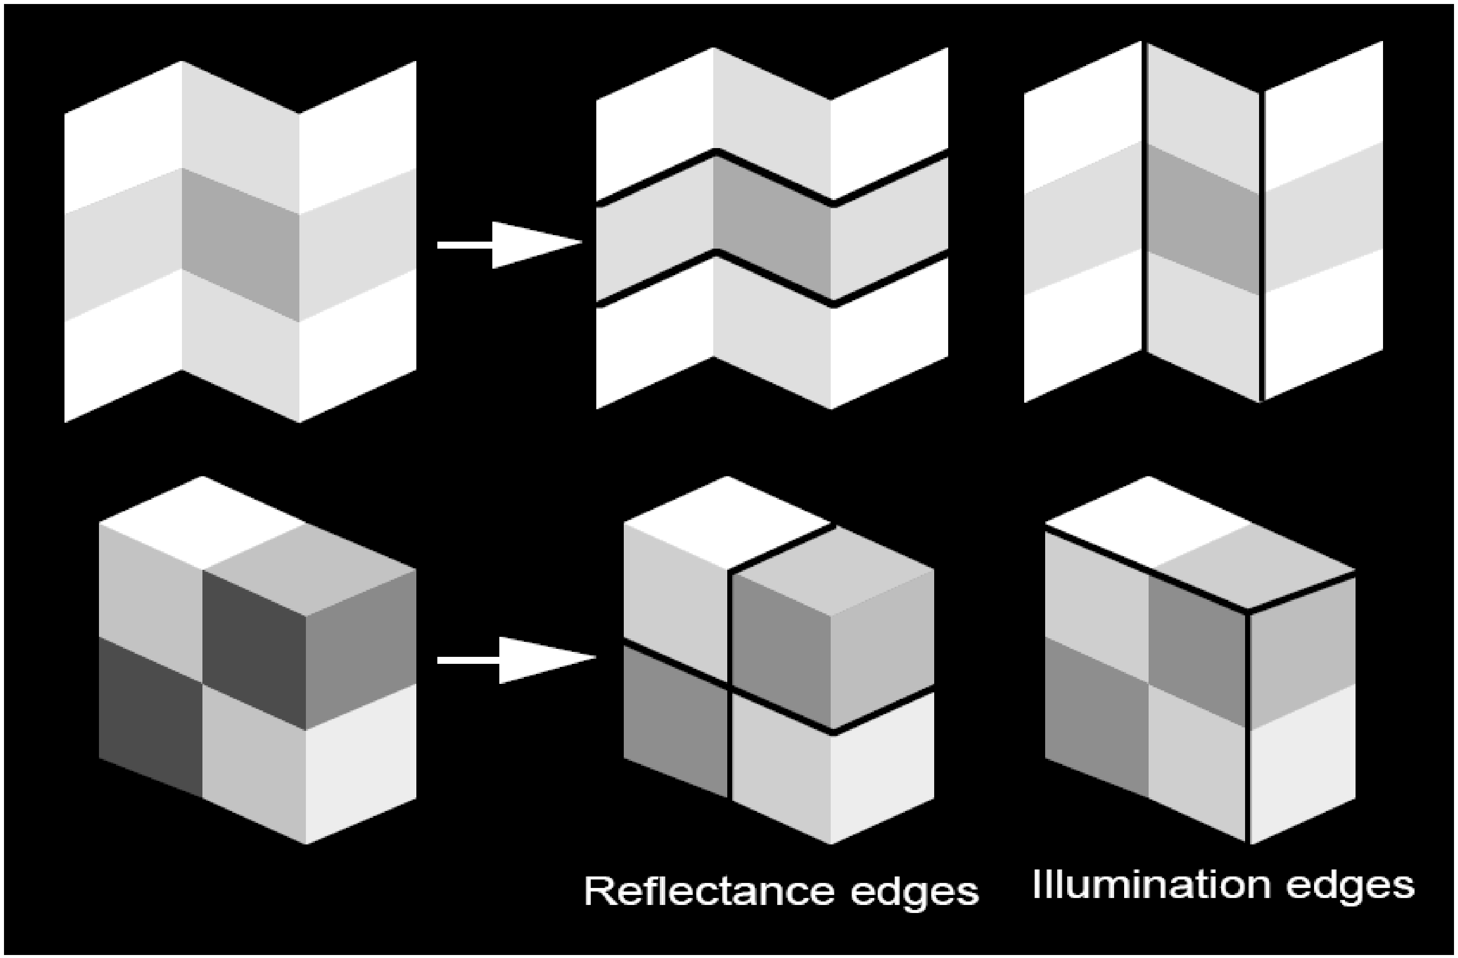

Examples

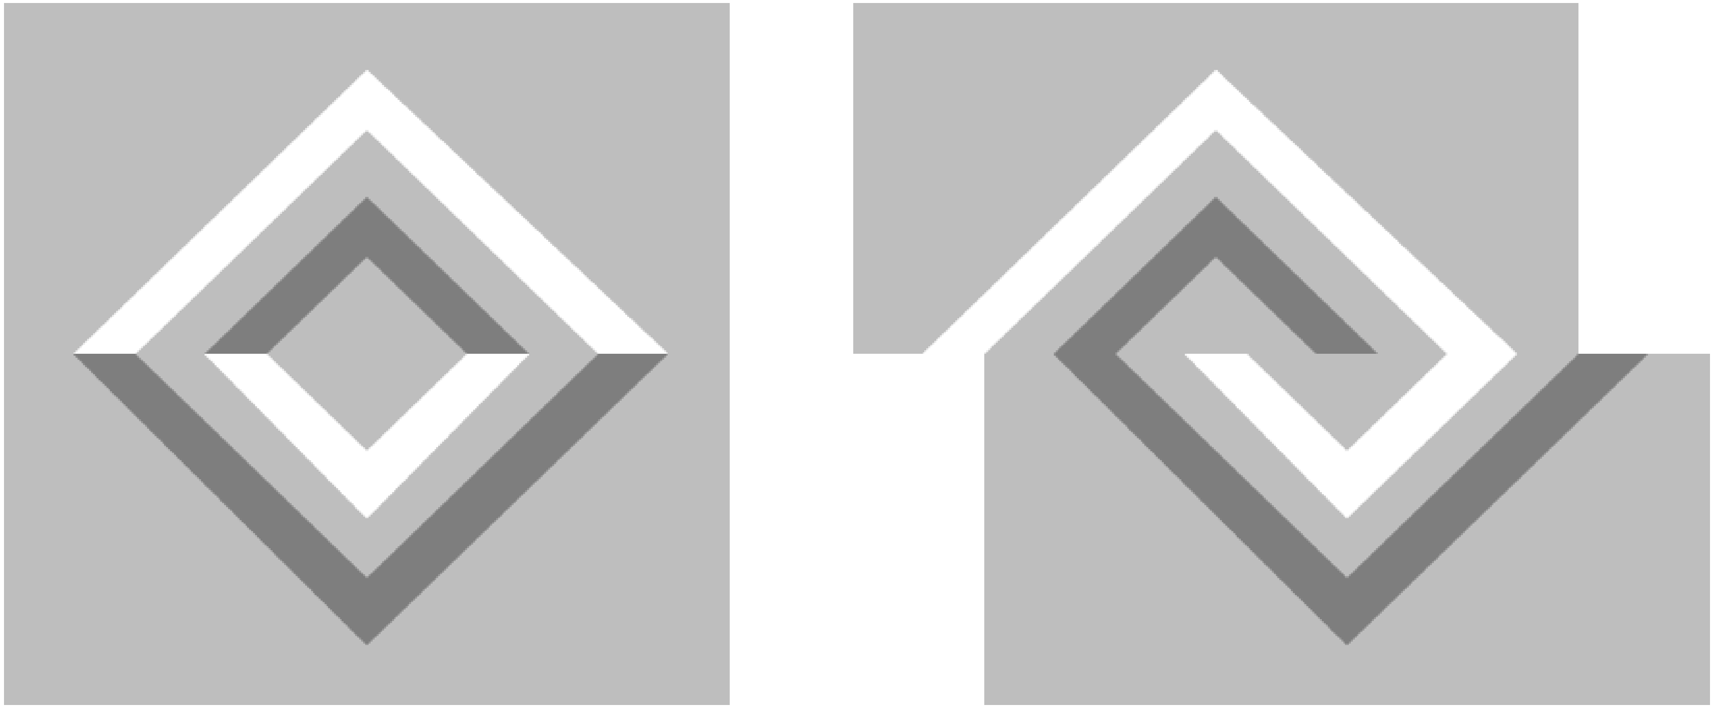

Junction Analysis of the 'Impossible' Object

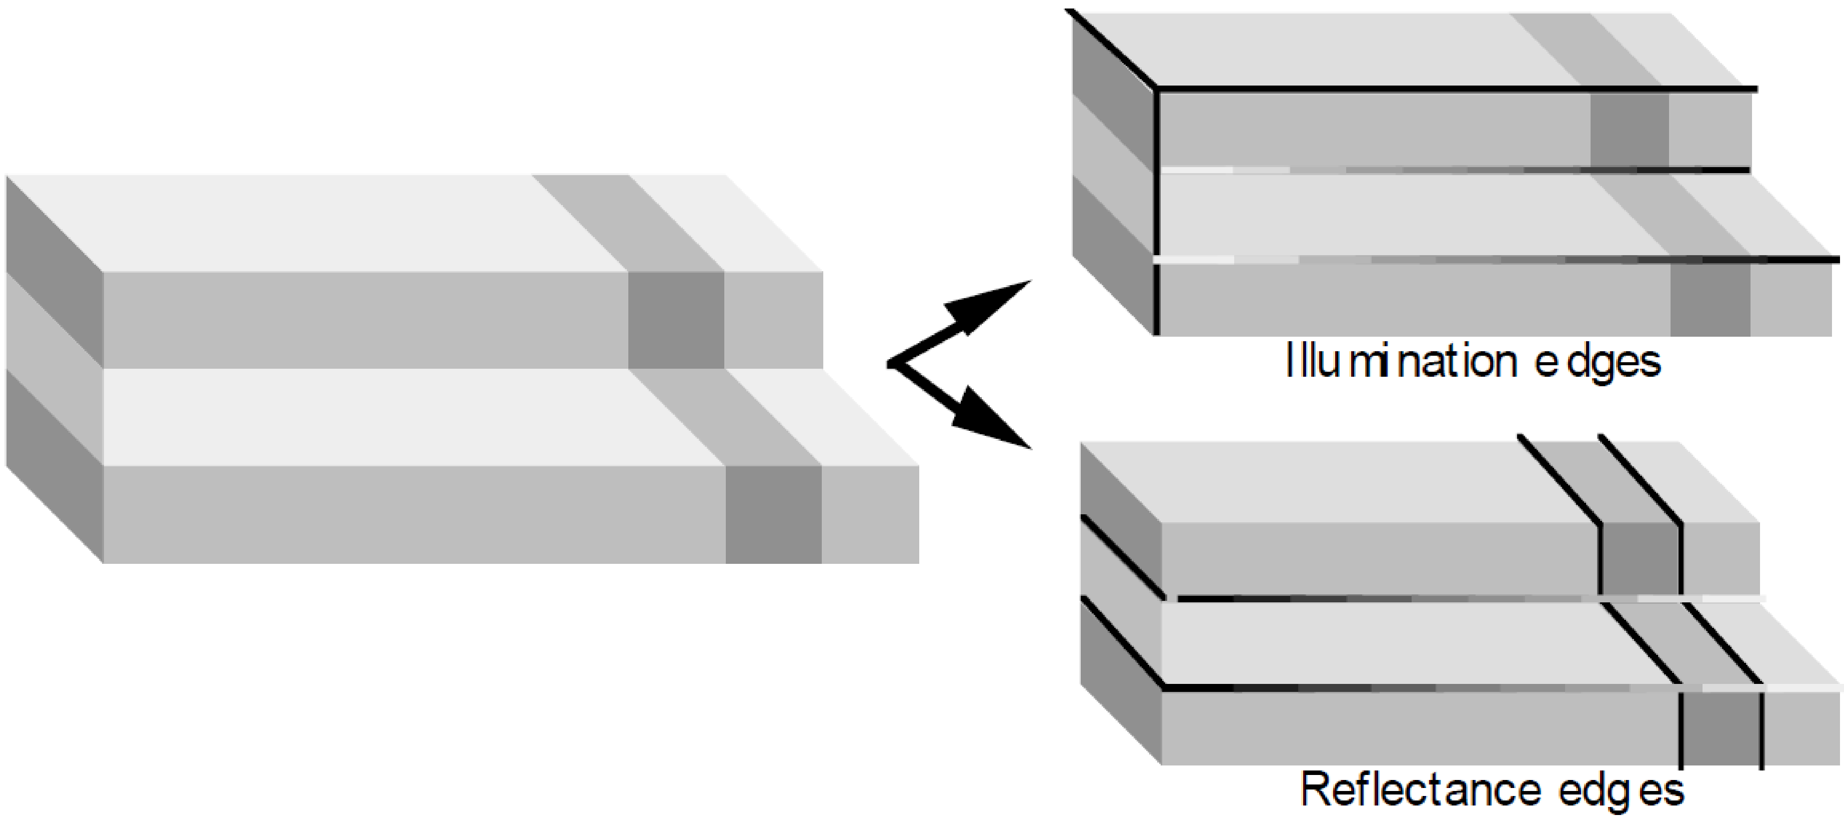

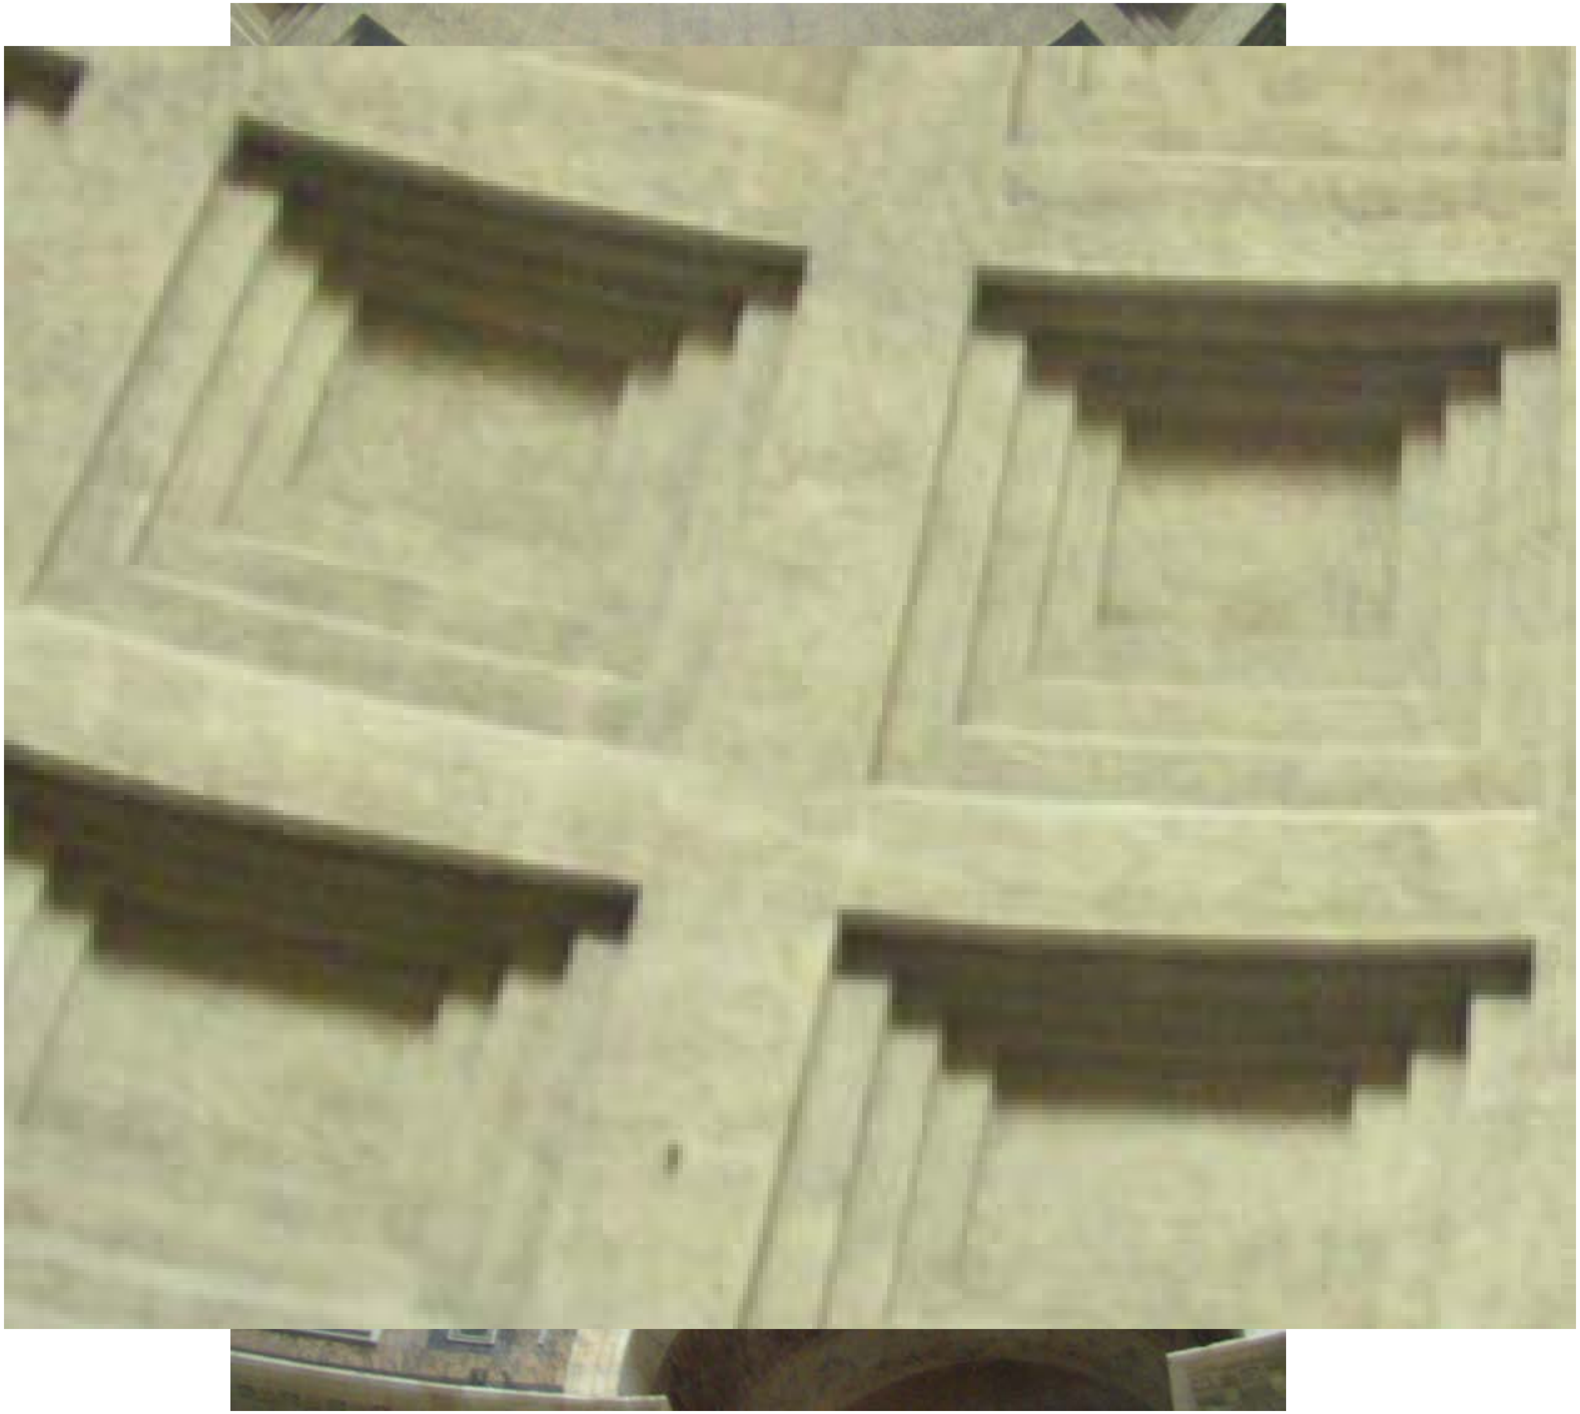

Counter-Example

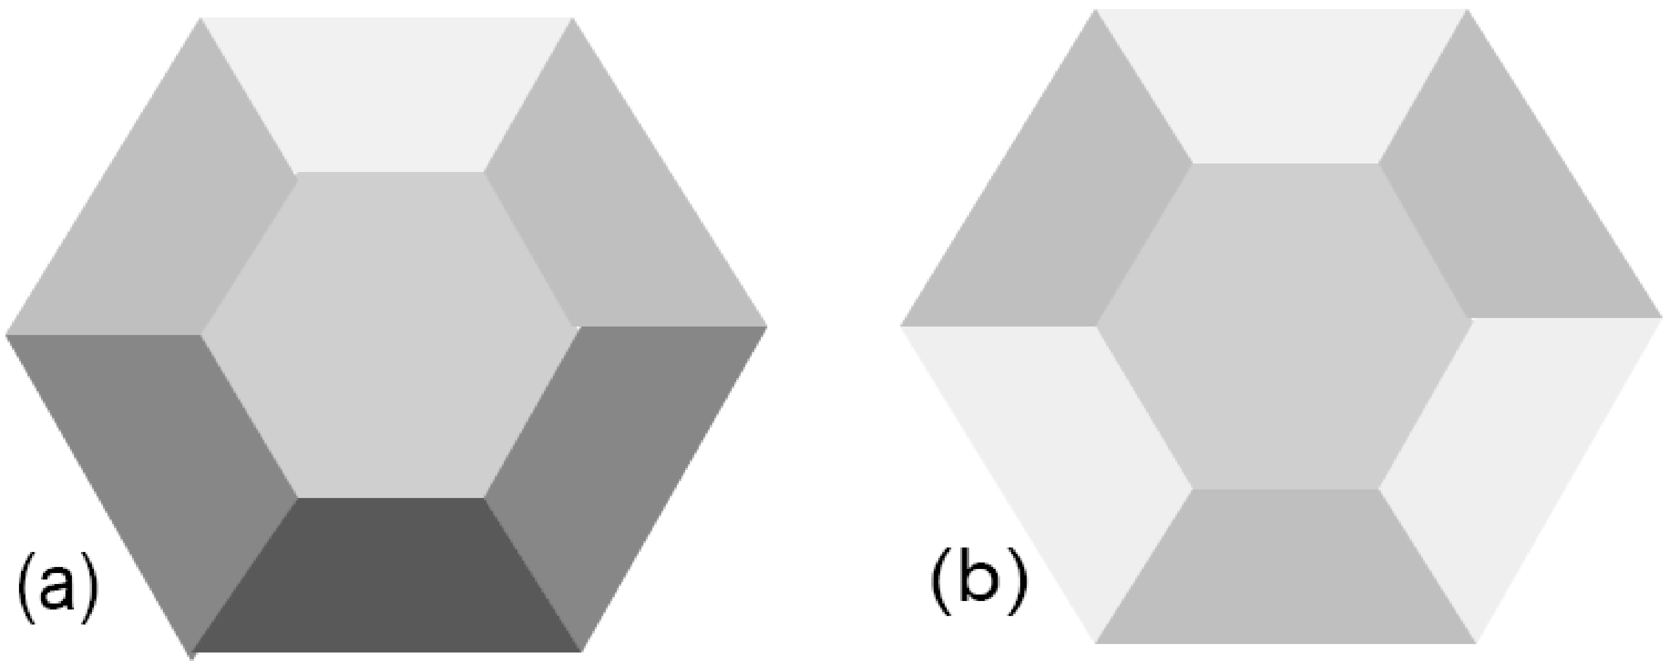

Consistency Check

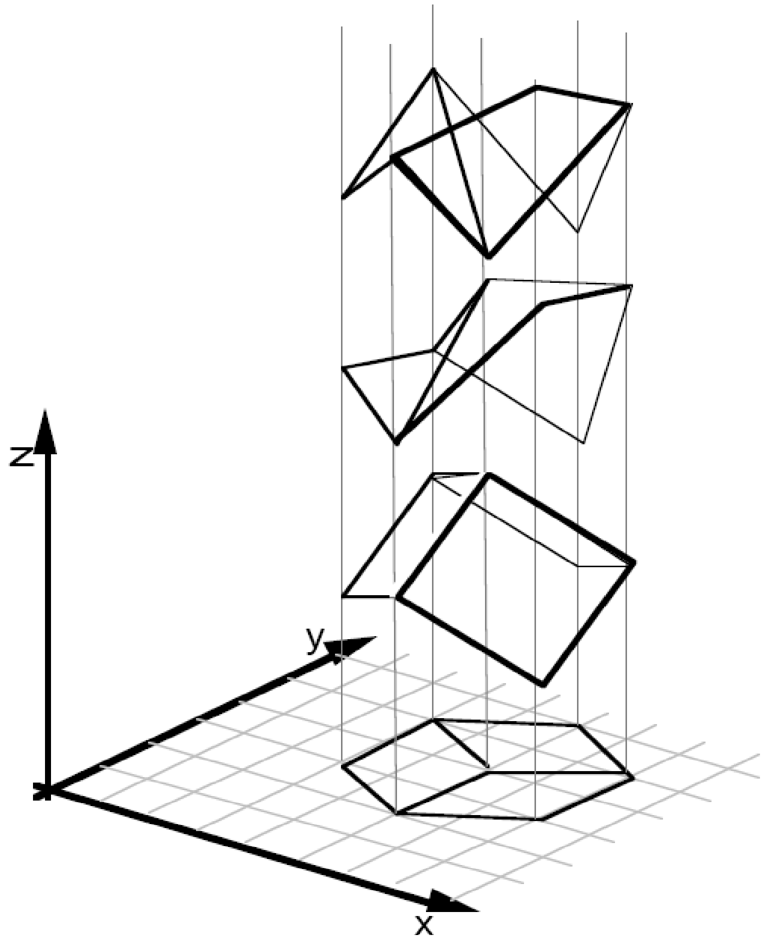

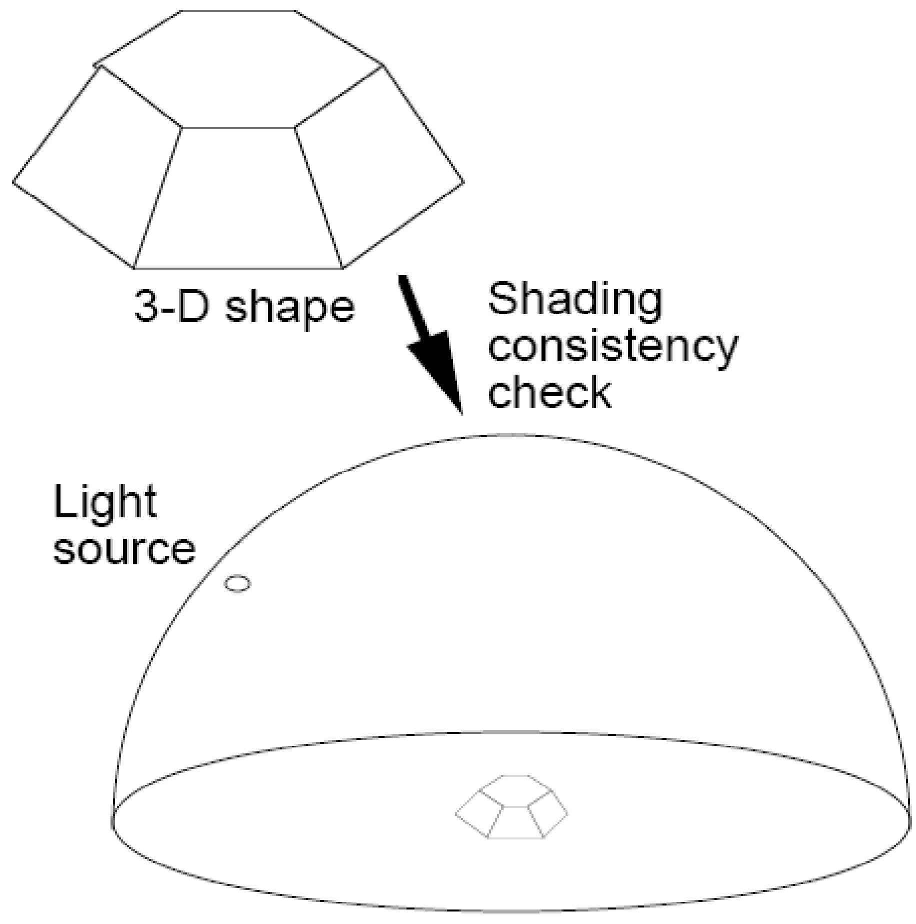

Global Measures of ‘Correctness’

- Low variance of angles
- Planarity of faces
- Overall compactness
- Consistency with light source

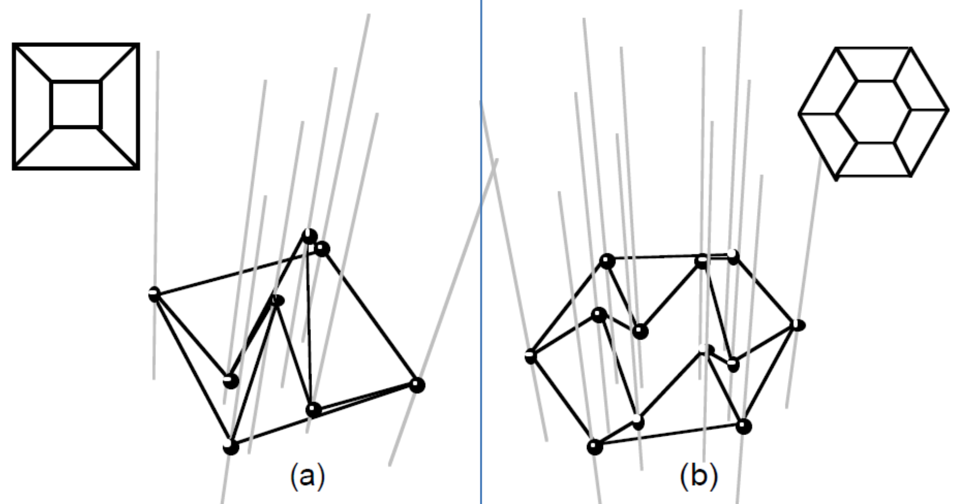

Global Measures of 'Correctness'

- Low variance of angles
- Planarity of faces
- Overall compactness
- Consistency with light source

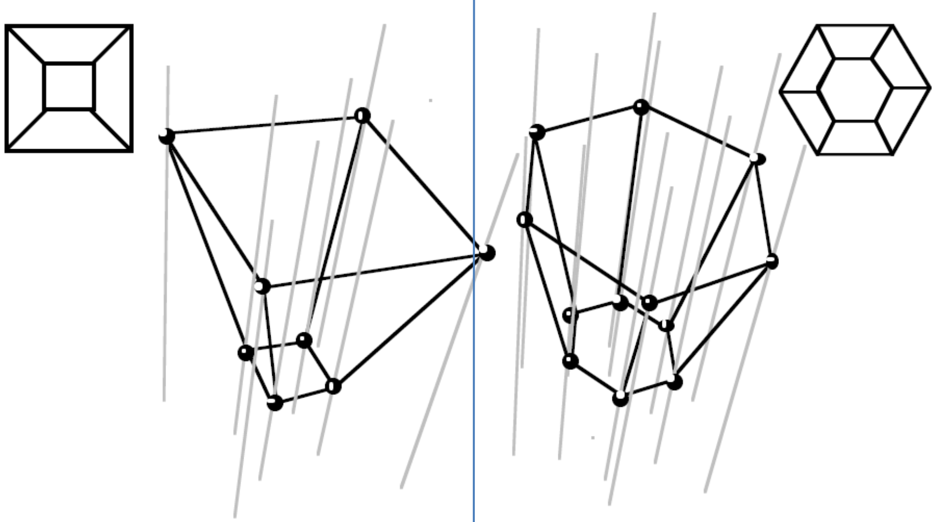

Possibility of Consistent Lighting

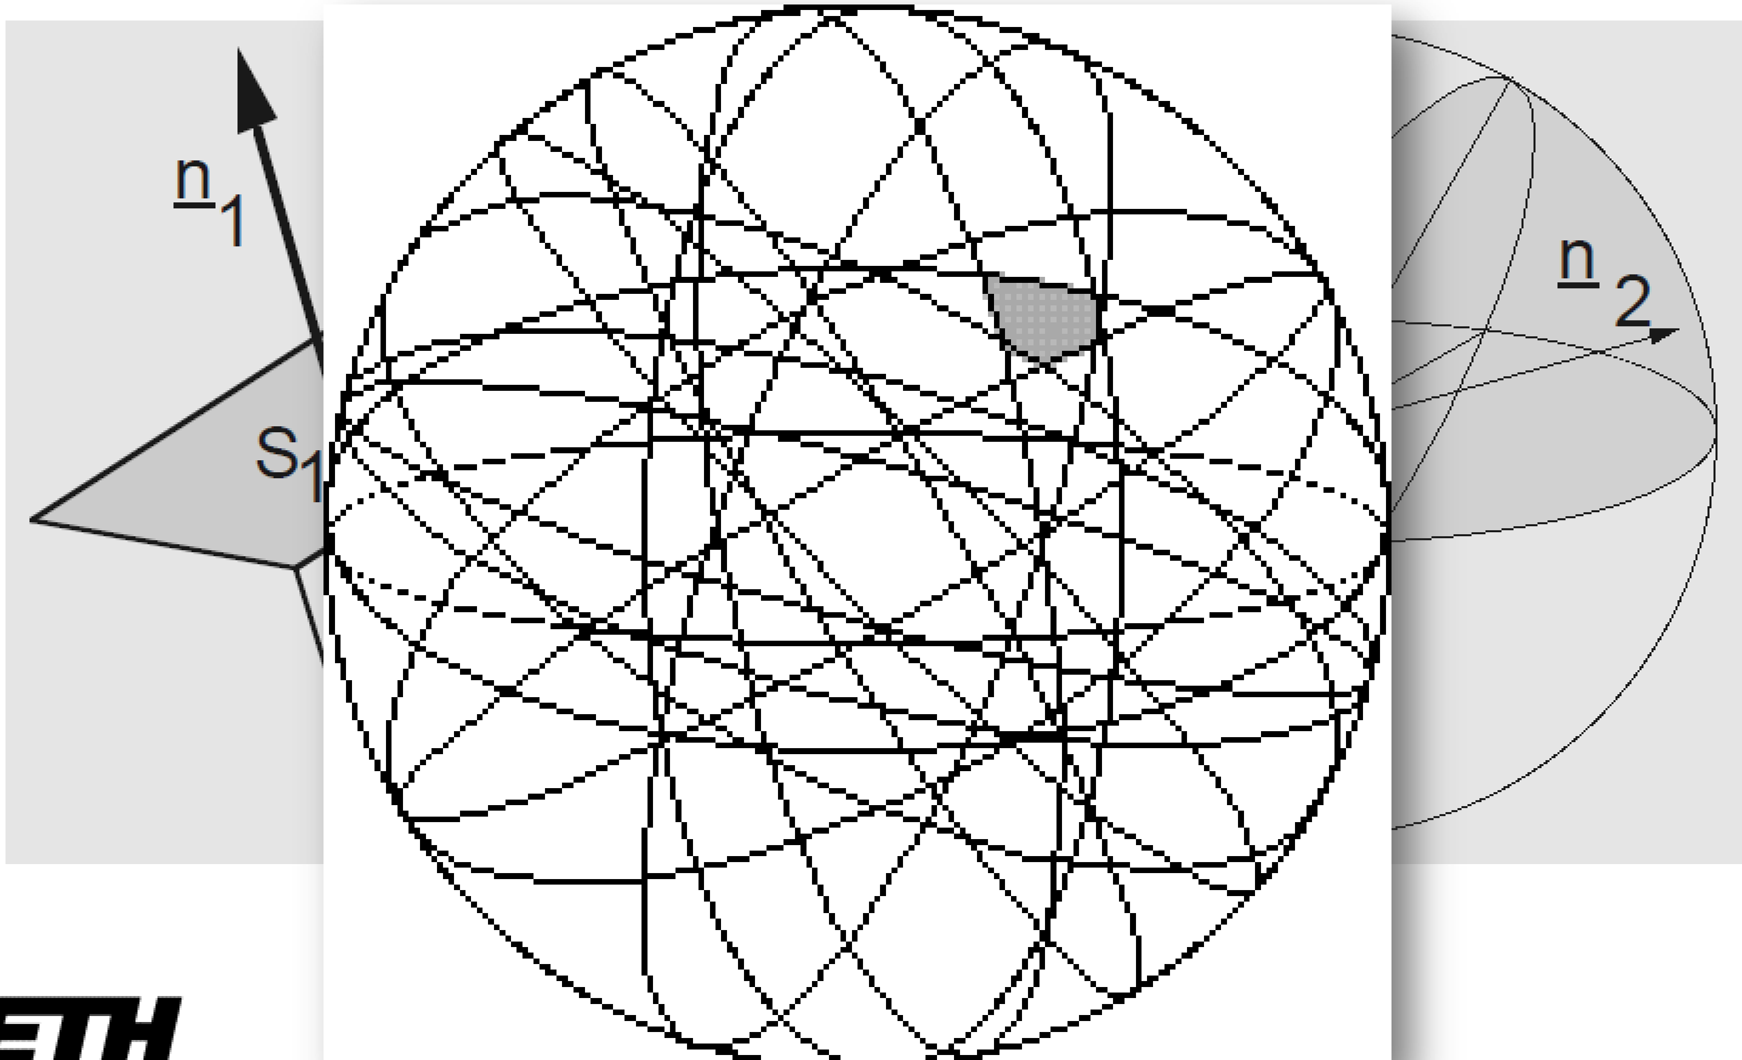

Global Analysis Confirms Local Analysis

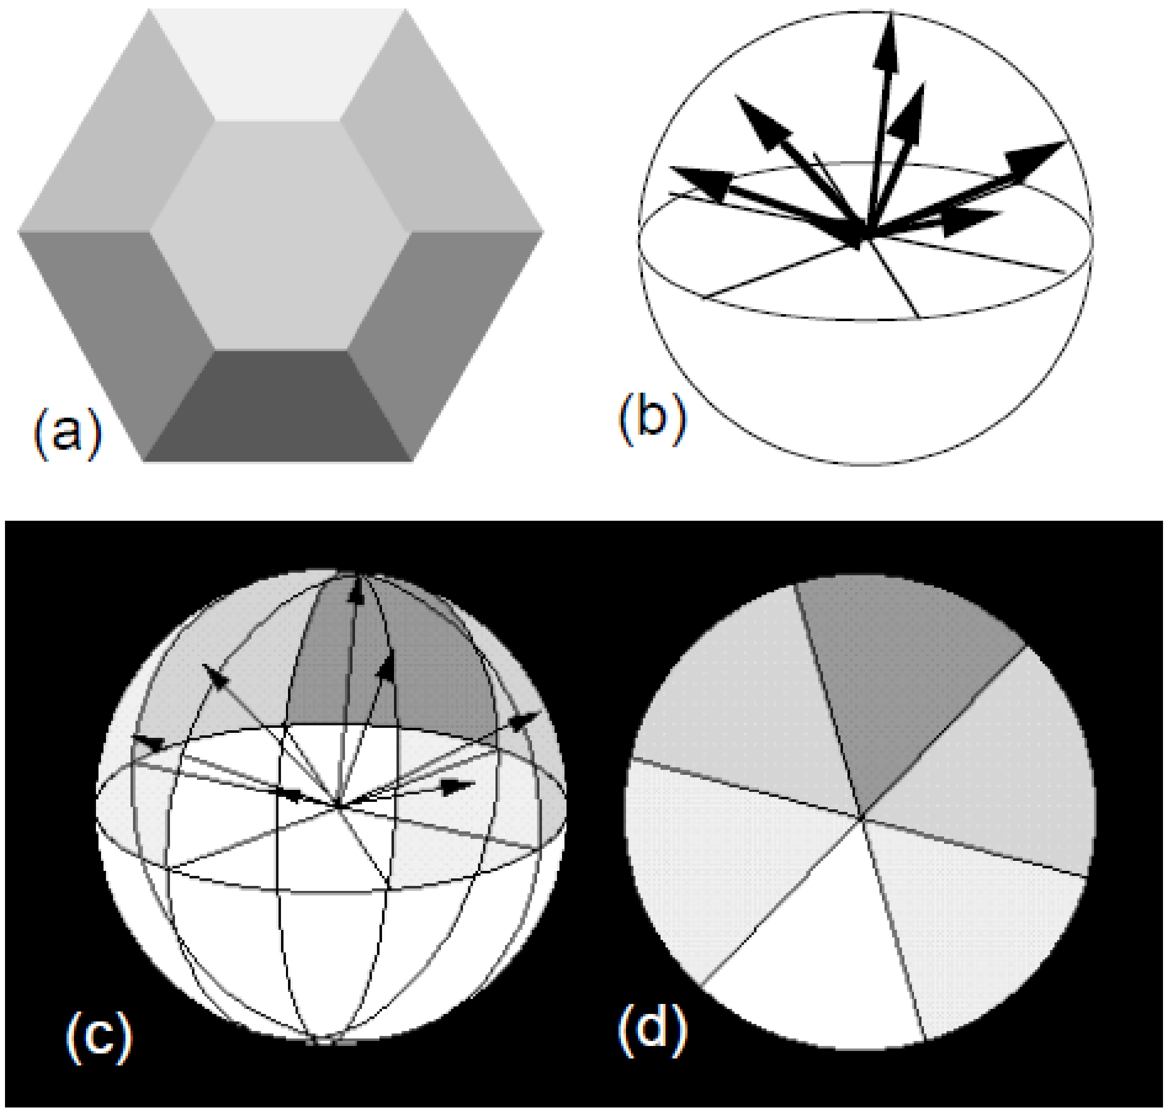

ETH

Global Analysis Trumps Local Analysis

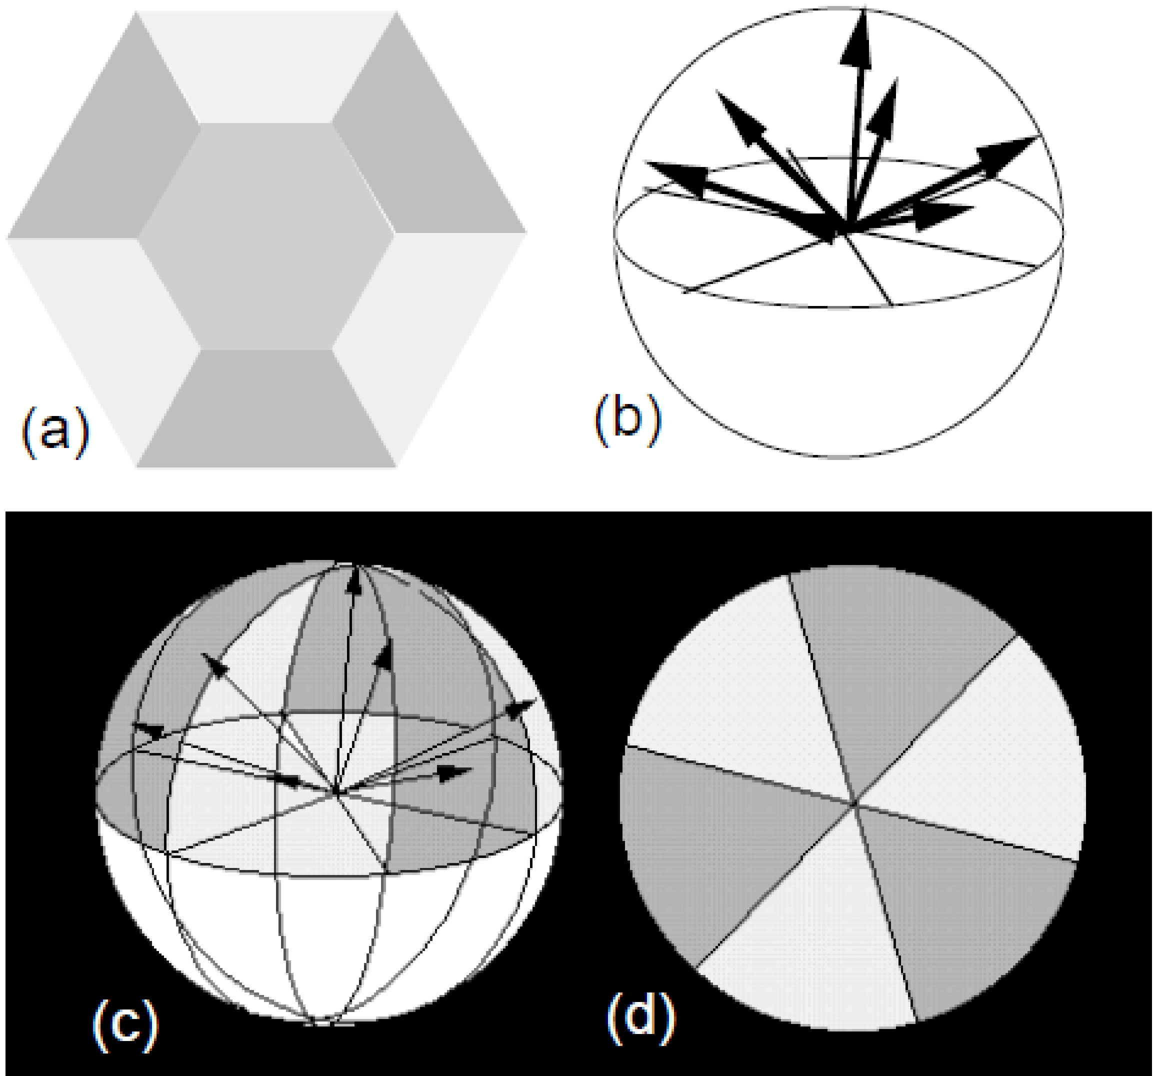

II

Image Sequences

- Deriving Intrinsic Images from Image Sequences
  - Weiss ICCV’01
- For static objects, multiple frames

a b c

Problem Formulation

Given a sequence of T images $\{I(x,y,t)\}_{t=1}^{T}$ in which reflectance is constant over time and only the illumination changes, can we solve for a single reflectance image and T Illumination images $\{L(x,y,t)\}_{t=1}^{T}$?

$$I(x,y) = L(x,y)R(x,y)$$

$$\Downarrow$$

$$\{I(x,y,t)\}_{t=1}^{T} = \{L(x,y,t)\}_{t=1}^{T} R(x,y)$$

Still completely ill-posed : at every pixel there are T equations and T+1 unknowns.

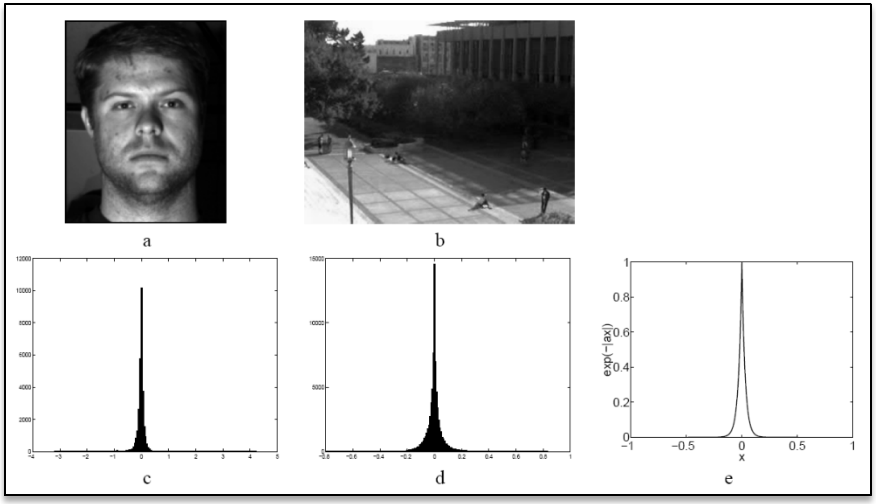

- Prior based on intuition:
  - derivative-like filter outputs of L tend to be sparse

$$\{I(x,y,t)\}_{t=1}^{T} = \{L(x,y,t)\}_{t=1}^{T} R(x,y)$$

(move to log-space)

$$i(x,y,t) = r(x,y) + l(x,y,t)$$

$$o_n(x,y,t) = i(x,y,t) * f_n$$

f_n = one of N filters like

1	-1

a b

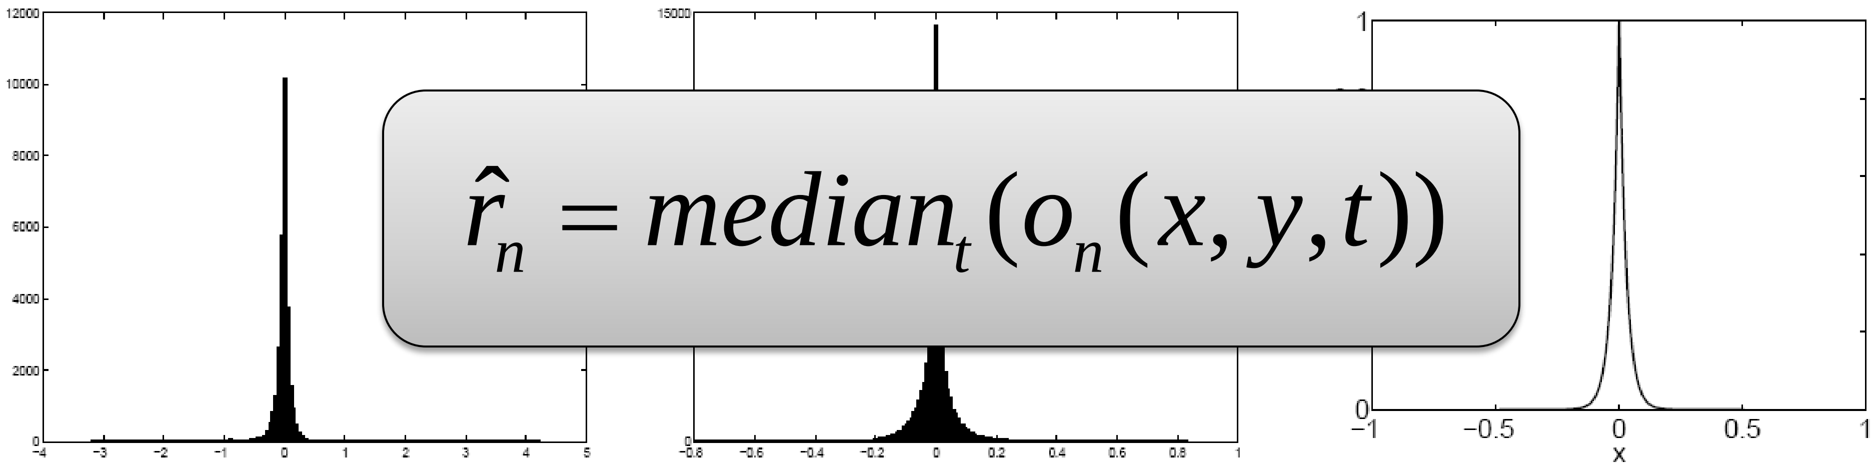

c d e

$$\hat{r}_n = median_t(o_n(x,y,t))$$

$$o_n(x,y,t) = i(x,y,t) * f_n$$

- Variety of responses has Laplacian-shaped distribution

Toy Example

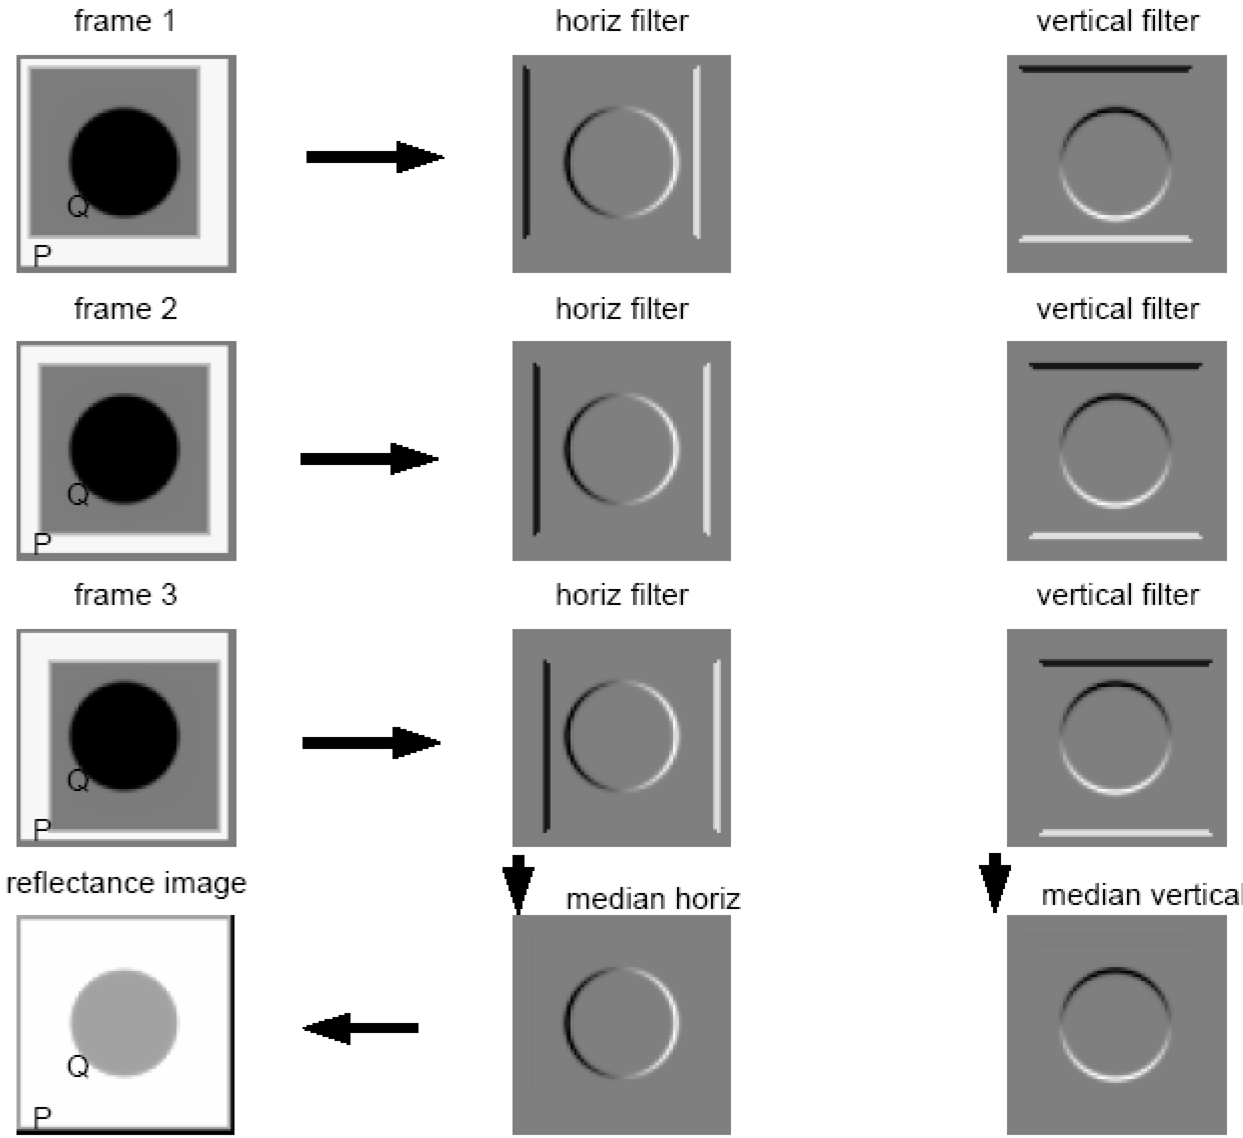
frame 1
horiz filter
vertical filter
Q
P
frame 2
horiz filter
vertical filter
Q
P
frame 3
horiz filter
vertical filter
Q
P
reflectance image
median horiz
median vertical
Q
P

Example Result 1

- Einstein image is translated diagonally 4 pixels per frame

Reagan image

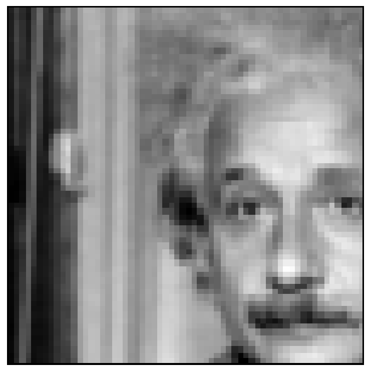

Einstein image

first frame

last frame

ML Reagan

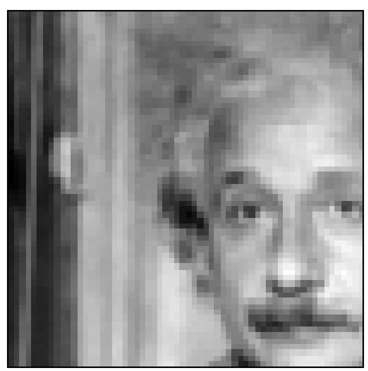

ML Einstein

min filter

median filter

Example Result 2

- 64 images with variable lighting from Yale Face Database


frame 2


frame 11


ML reflectance


ML illumination 2


ML illumination 11

III

Single Image

- Recovering Intrinsic Images from a Single Image
  - Tappen, Freeman, Adelson
    - NIPS’03 & PAMI’05

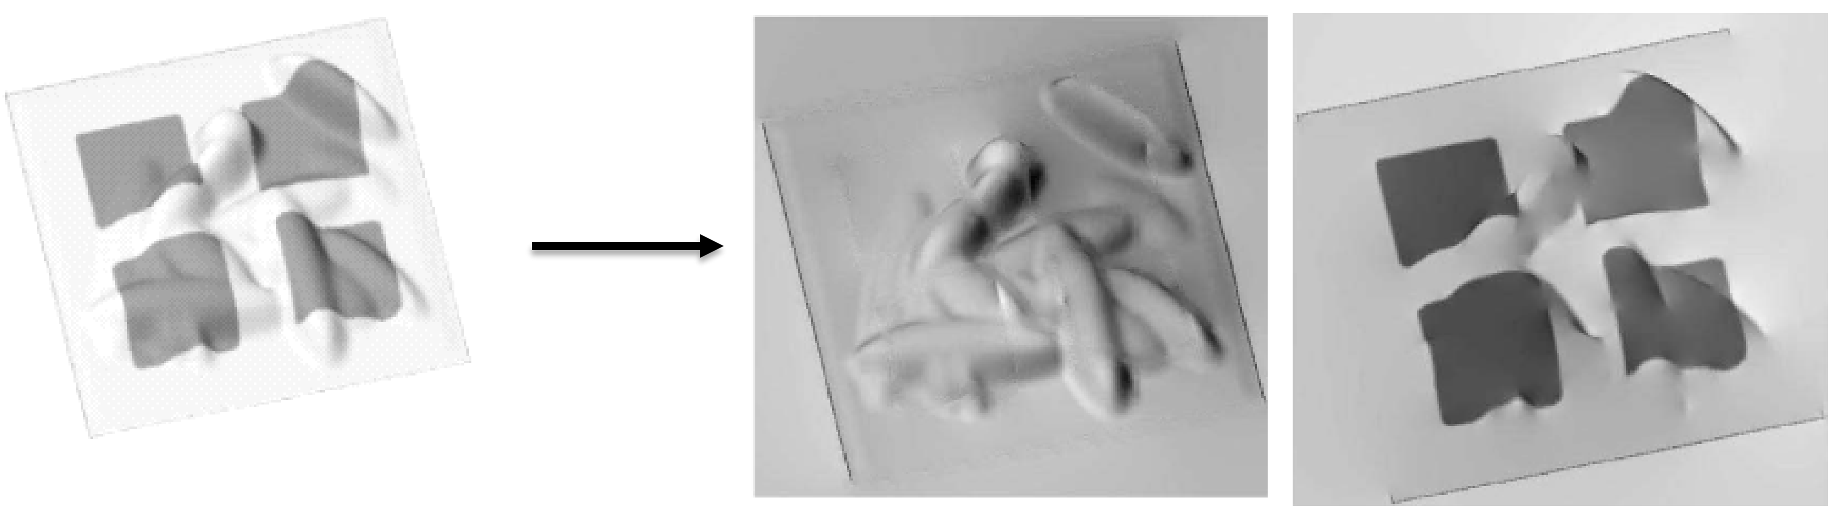

ETH

Assumption

- Each derivative is caused either by Shading or Reflectance
- Reduces to a binary classification problem

Image Derivative w.r.t. x and y

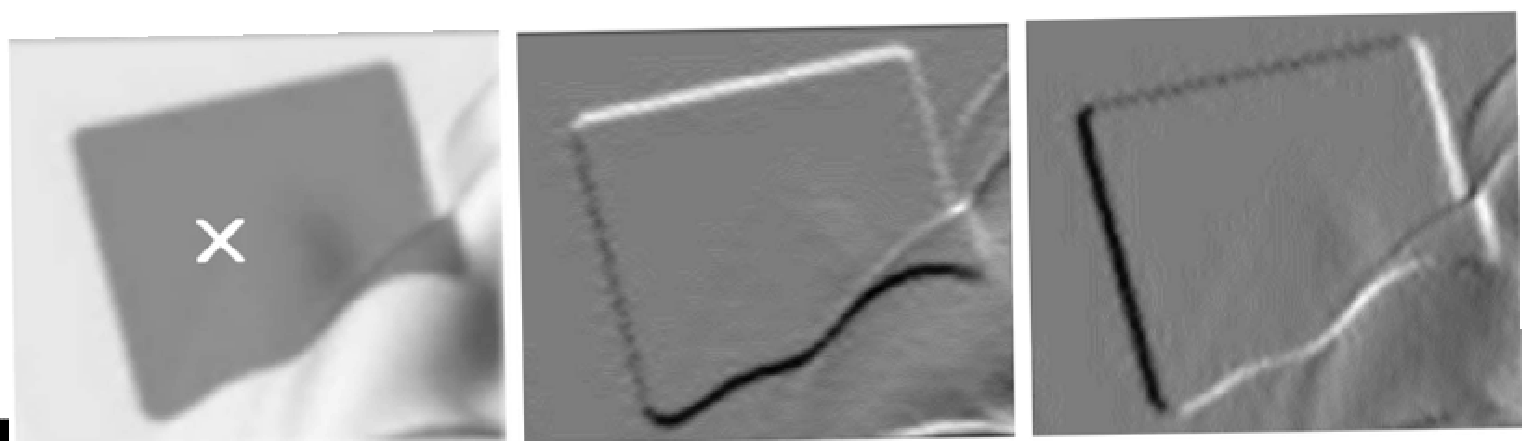

Classifying Derivatives

- 4 Basic phases:
  1. Compute image derivatives
  2. Classify each derivative as caused by shading or reflectance
  3. Invert derivatives classified as shading to find shading images
  4. Reflectance image is found the same way

Classification

1. Color information

- changes due to shading should affect R,G and B proportionally

$$C_1 = \alpha \cdot C_2$$

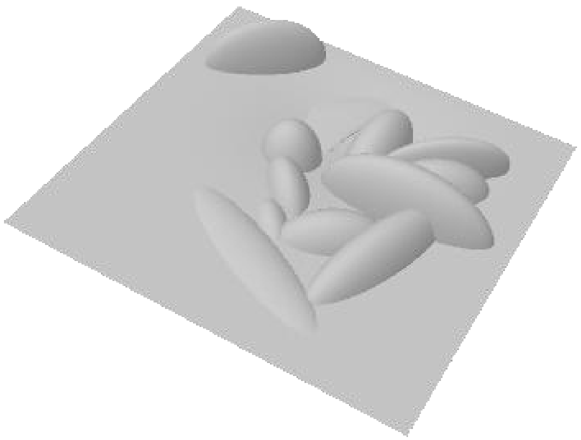

If $C_1 \neq \alpha \cdot C_2$ the changes are caused by reflectance

Color Information - examples

Black on white may be interpreted as intensity change.

Resulting in misclassification

Classification

1. Color information

- changes due to shading should affect R,G and B proportionally

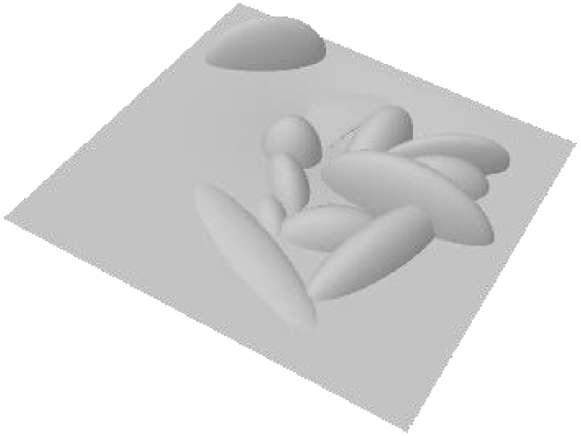

$$C_1 = \alpha \cdot C_2$$

If $C_1 \neq \alpha \cdot C_2$ the changes are caused by reflectance

Classification

1. Color information

- changes due to shading should affect R,G and B proportionally

$$C_1 = \alpha \cdot C_2$$

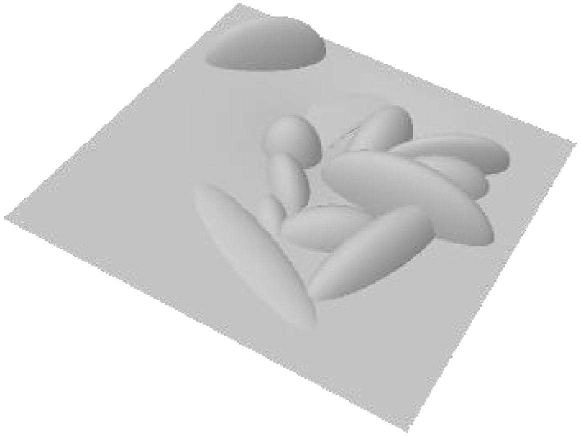

If $C_1 \neq \alpha \cdot C_2$ the changes are caused by reflectance

2. Statistical regularities of surfaces

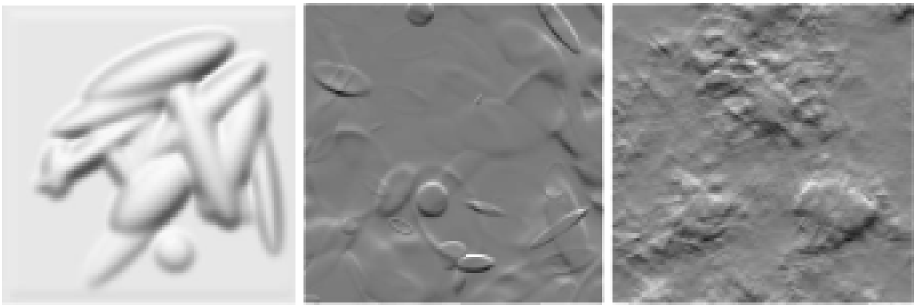

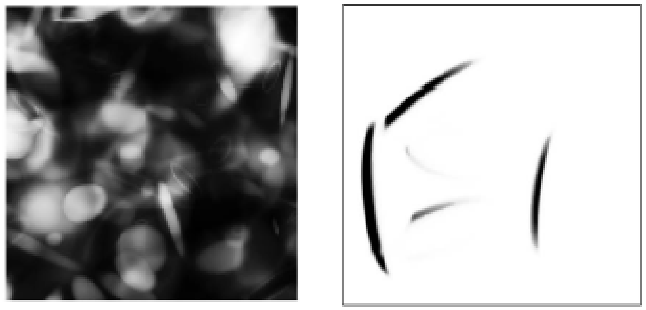

GrayScale Information - examples

Misclassification of the cheeks – due to weak gradients

Combing Information (Assuming Statistical Indep.)

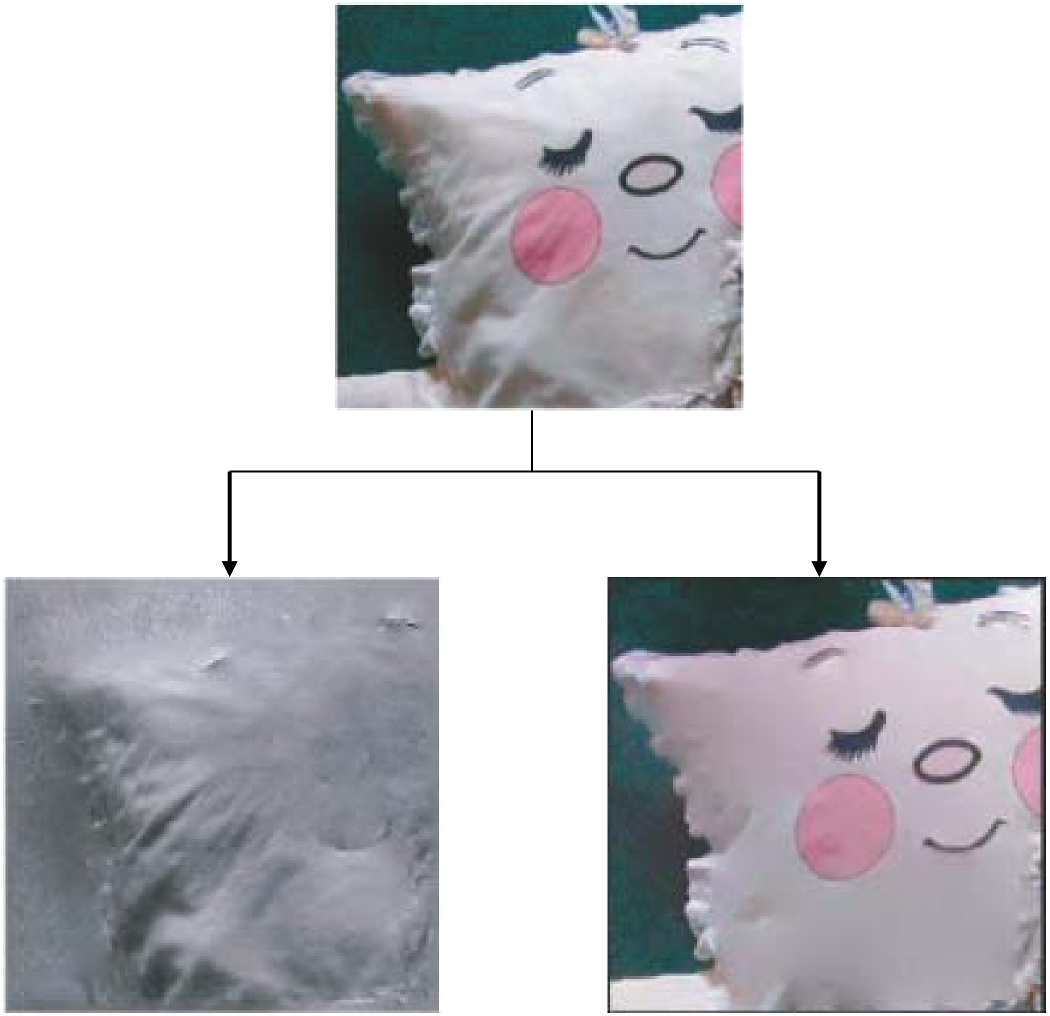

Handling Ambiguities

- Ambiguities - for example – center of the mouth

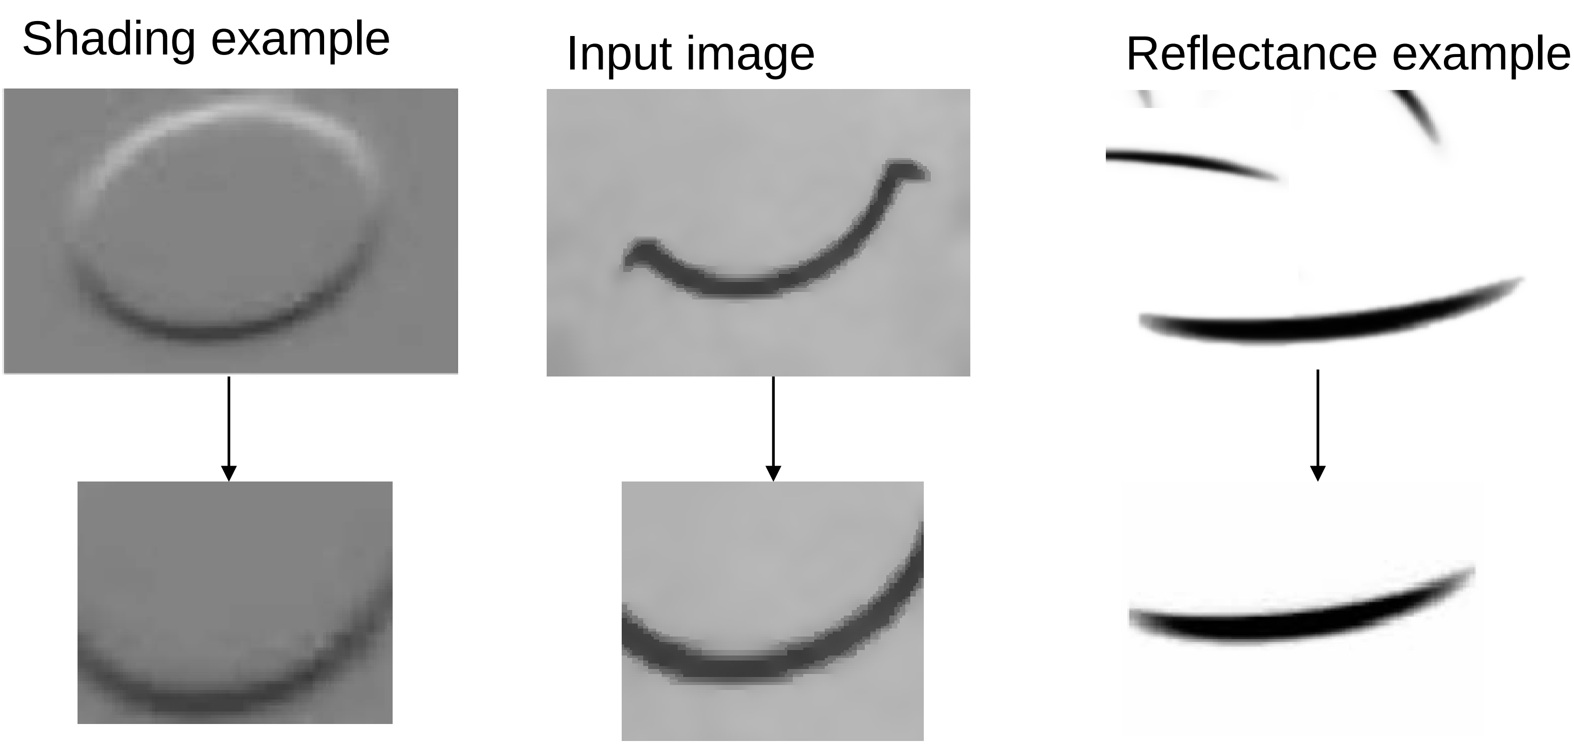

Handling Ambiguities

- Derivatives that lie on the same contour should have the same classification

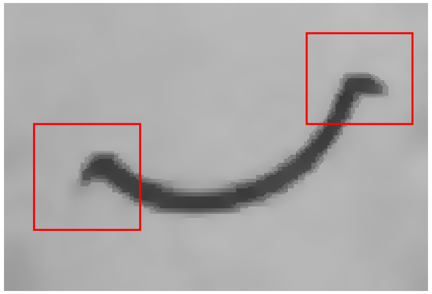

- The mouth corners are well classified as reflectance

→ Propagate evidence from conclusive areas to ambiguous ones using MRF

Final Results

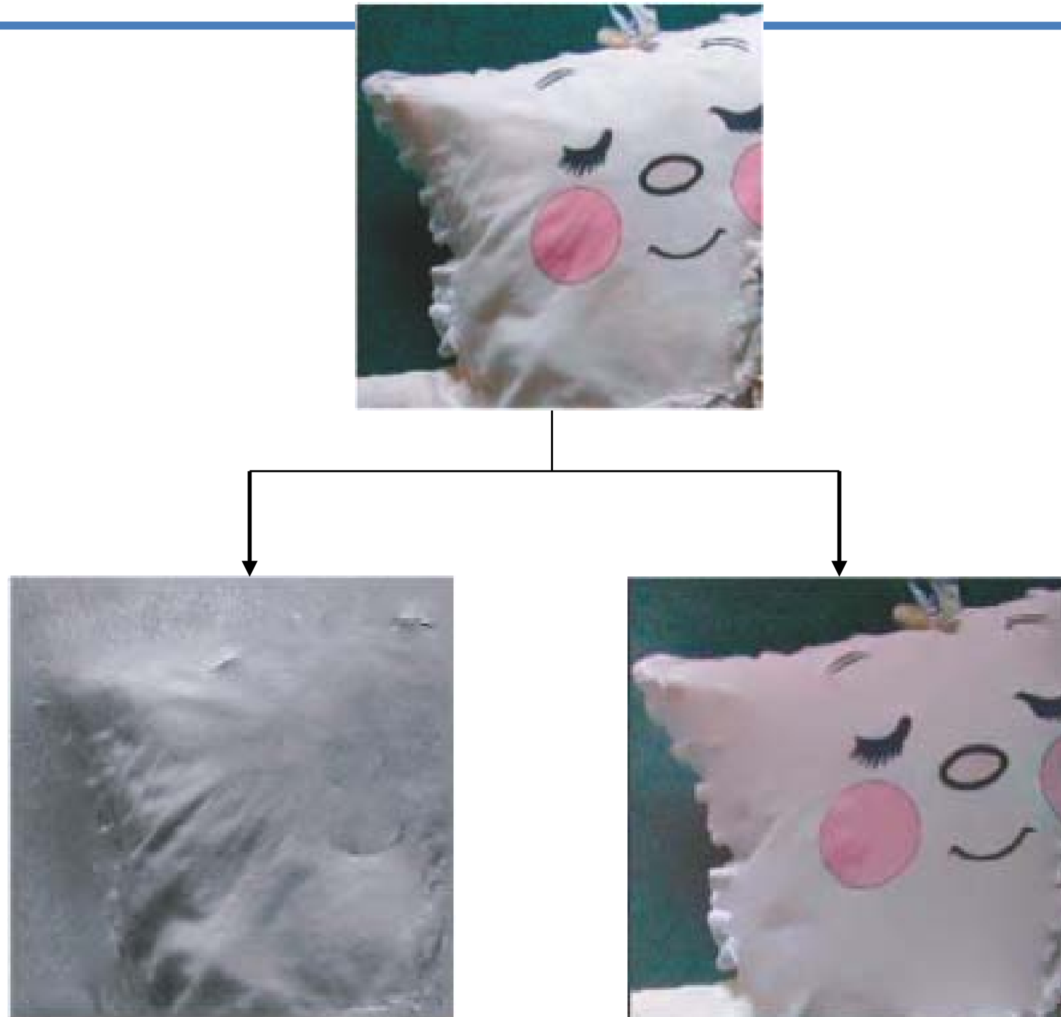

Unpropagated

Final Results

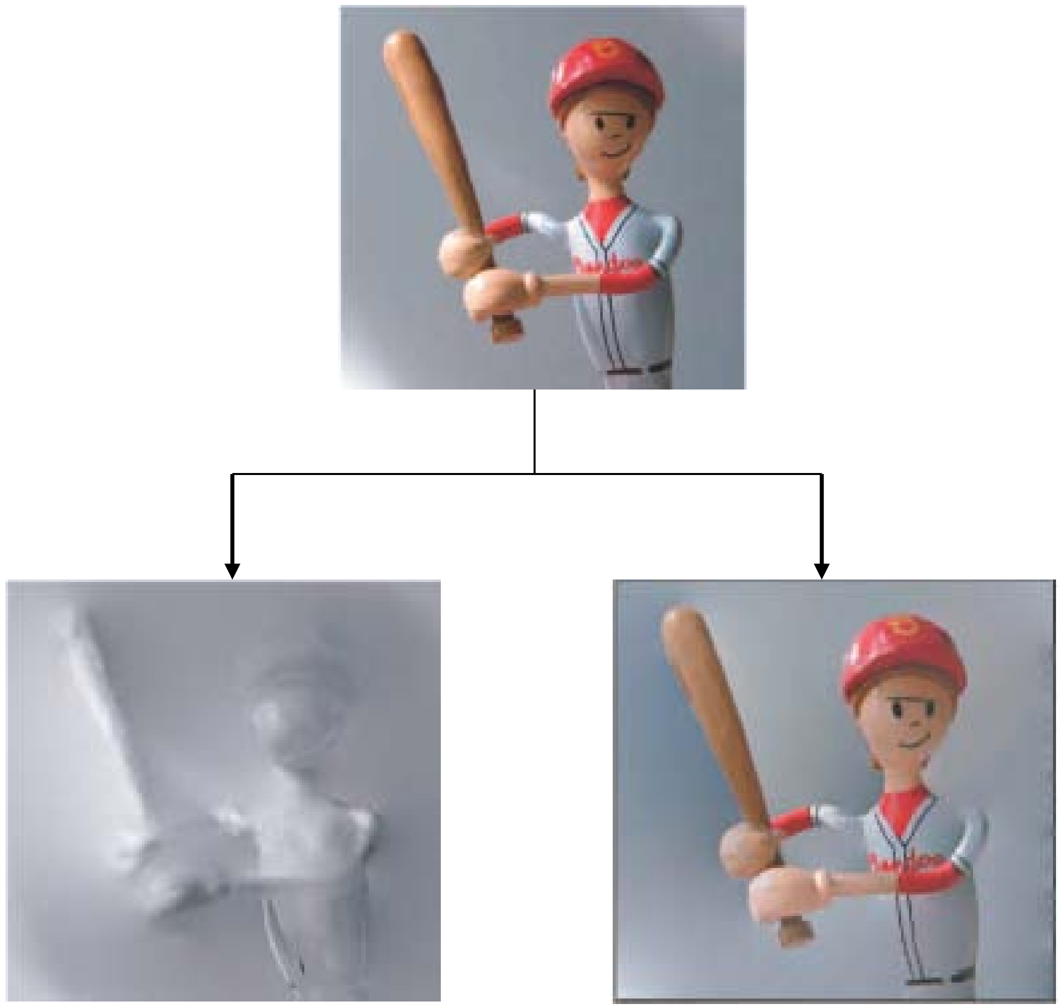

Final Results

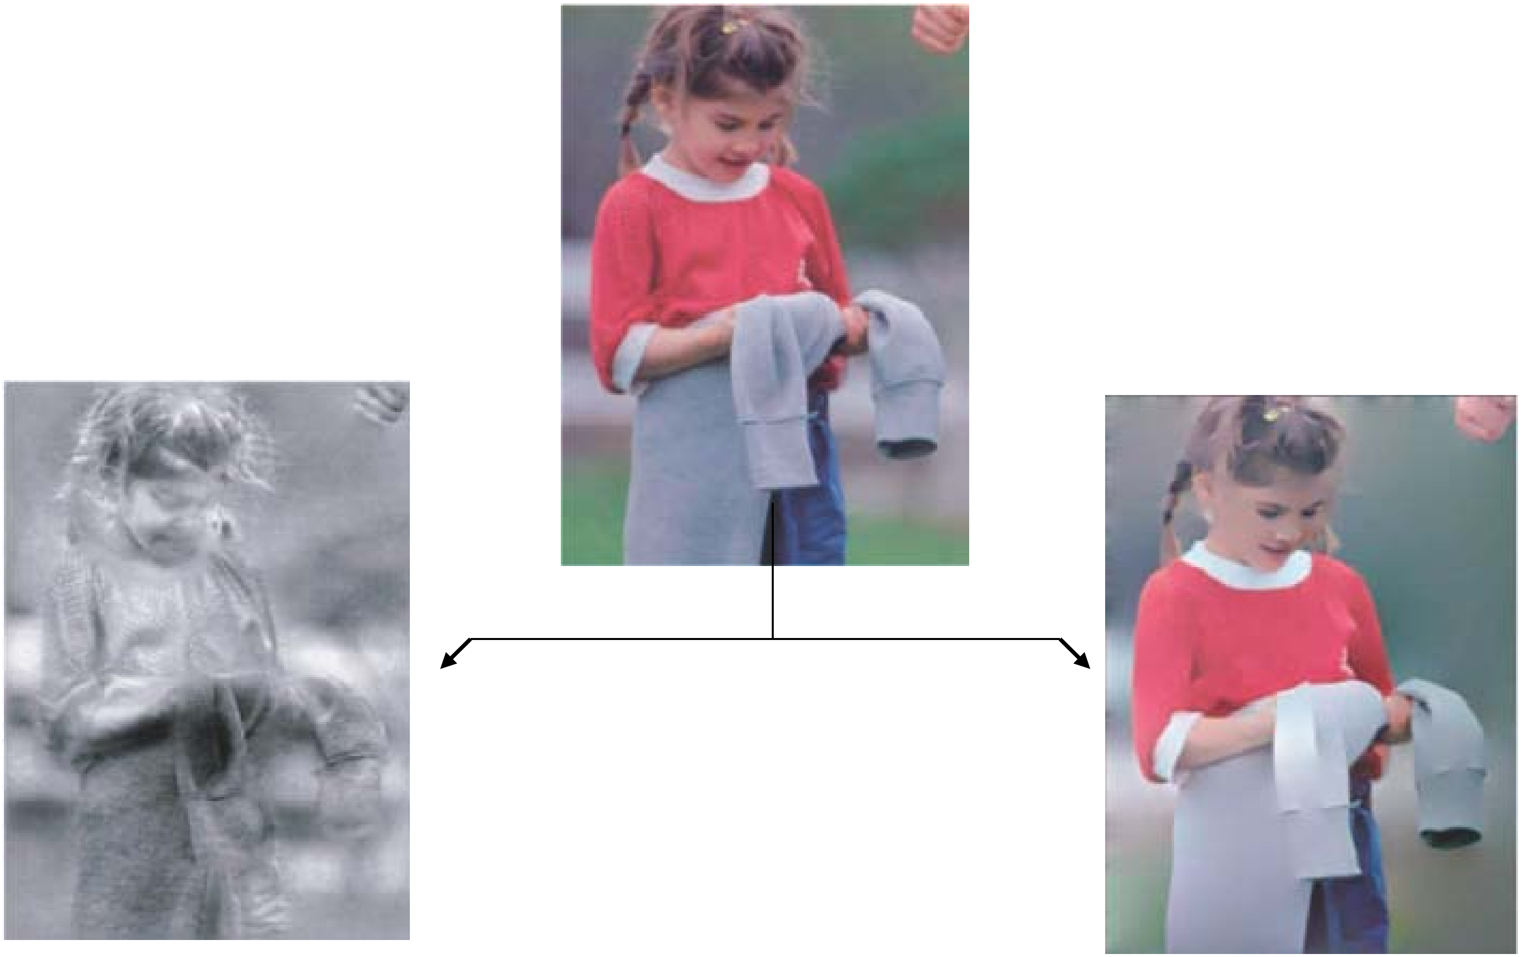

Final Results

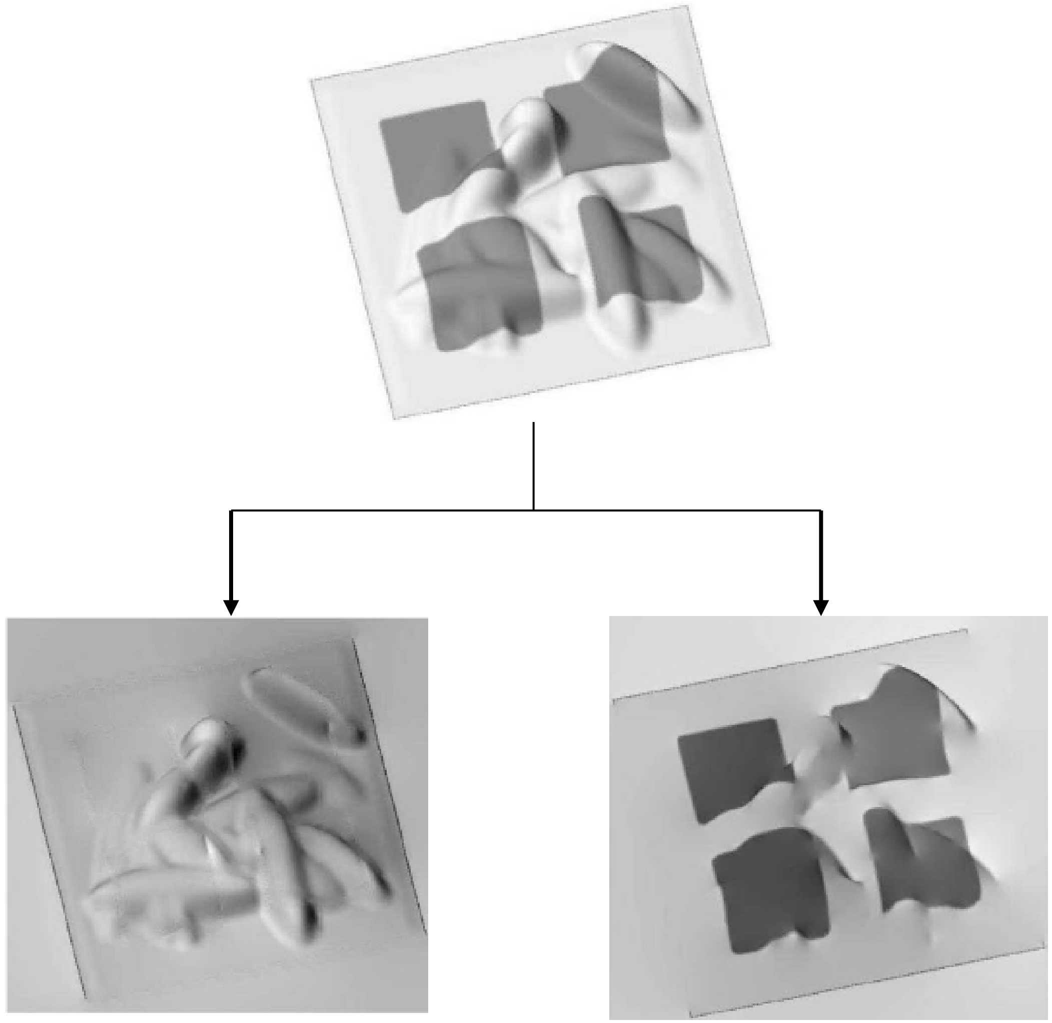

IV

Entropy Minimization

- Intrinsic Images by Entropy Minimization
  - Finlayson, Drew, Lu, ECCV'04

ETH

Sensor Response at a Pixel

$$\boxed{p_k = \int_{\lambda} R(\lambda) L(\lambda) S_k(\lambda) d\lambda}$$

R = Reflectance

L = Illumination

S = Sensor Sensitivity

Best When Sensors are Narrow Band

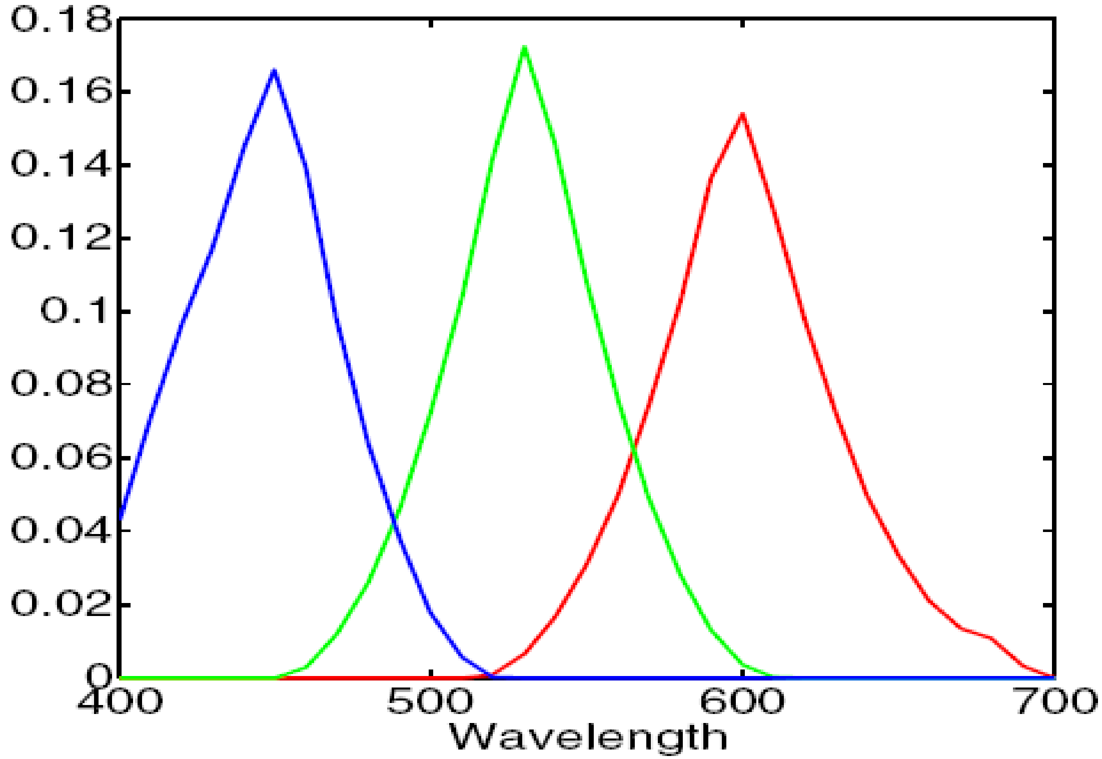

Best When Sensors are Narrow Band

$$S_k(\lambda) = \delta(\lambda - \lambda_k)$$

$$k \in \{R, G, B\}$$

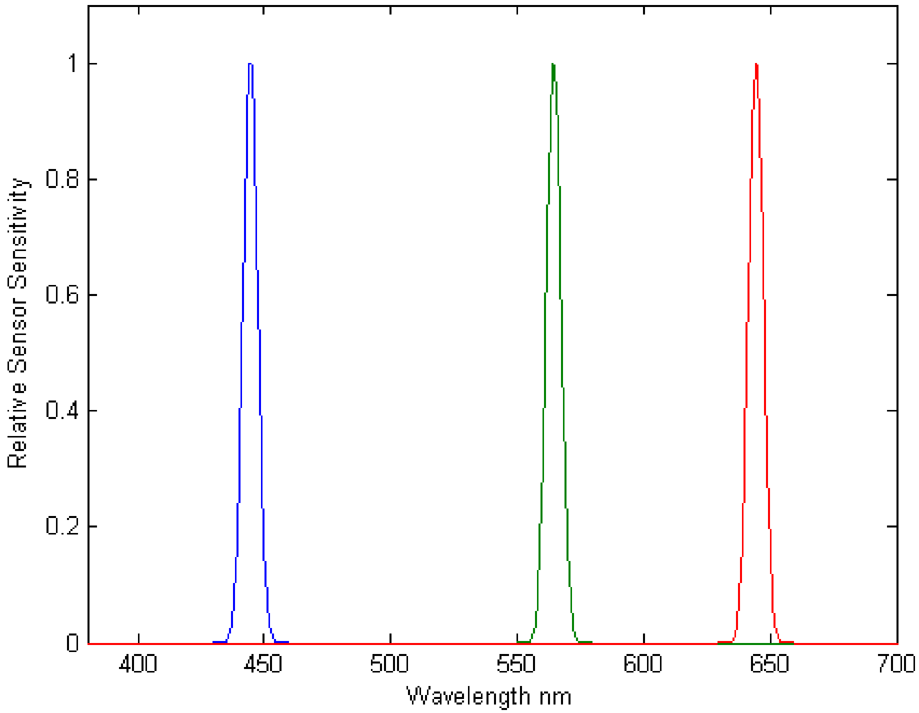

Just Reflectance & Illumination

$$p_k = \int_\lambda R(\lambda)L(\lambda)S_k(\lambda)d\lambda$$

$$\Downarrow$$

$$p_k = \int_\lambda R(\lambda)L(\lambda)\delta(\lambda - \lambda_k)d\lambda$$

$$\Downarrow$$

$$p_k = R(\lambda_k)L(\lambda_k)$$

Chromaticity for 7 Surfaces for 10 Illuminants

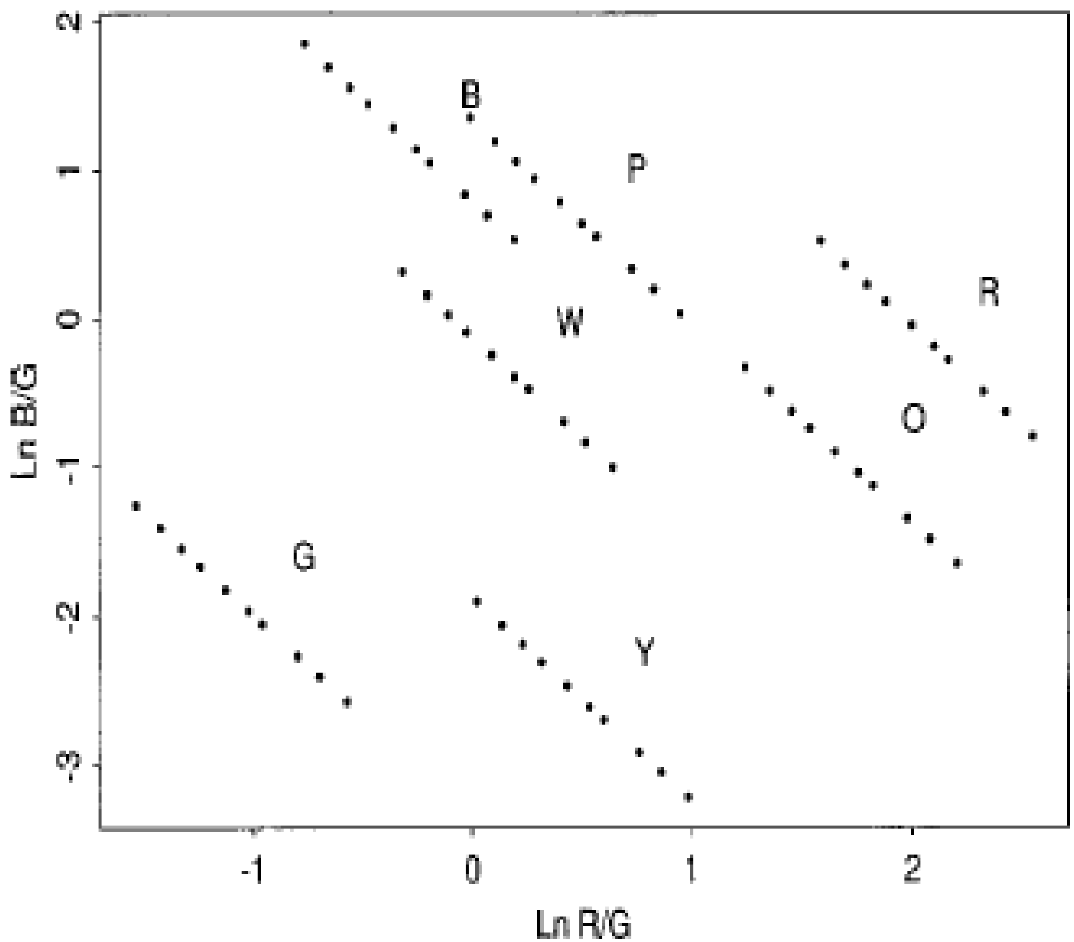

Macbeth Chart Under Changing Illumination

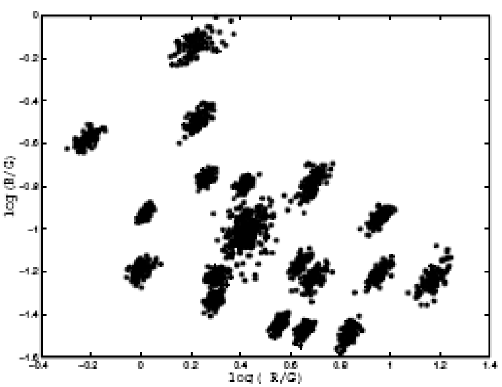

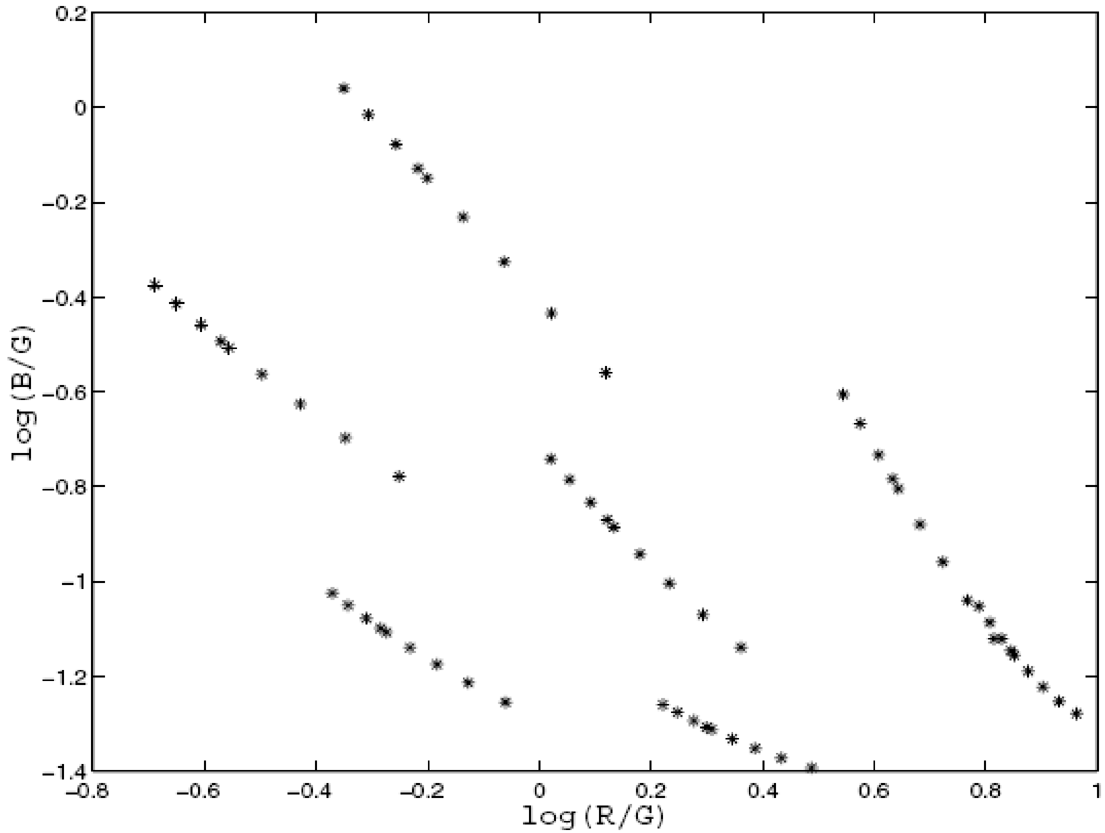

Entropy Minimization

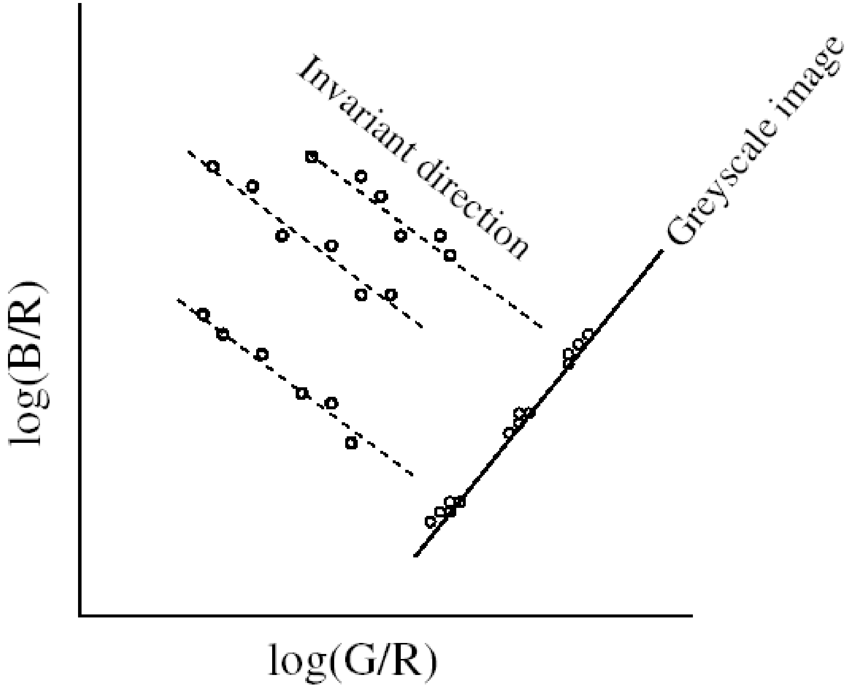

Correct Projection

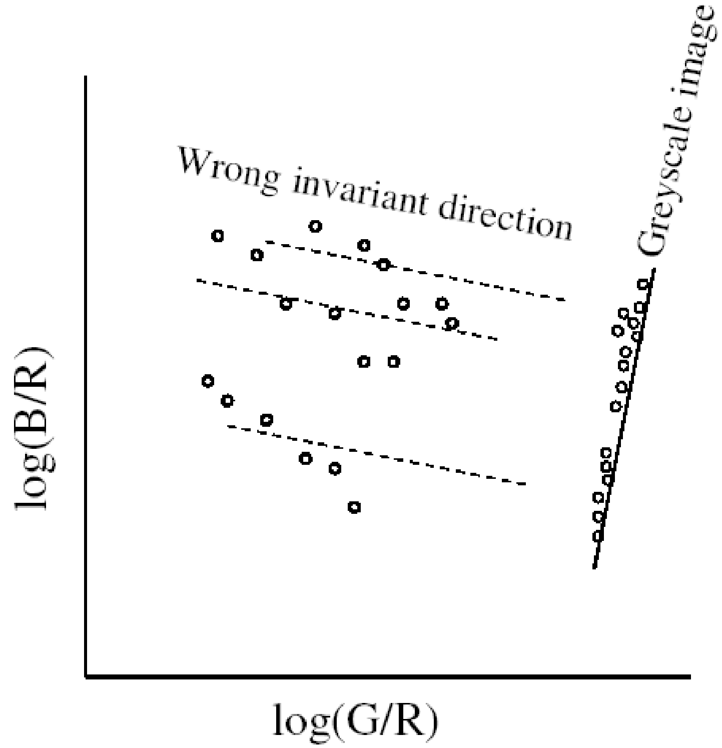

Incorrect Projection

Entropy Minimization

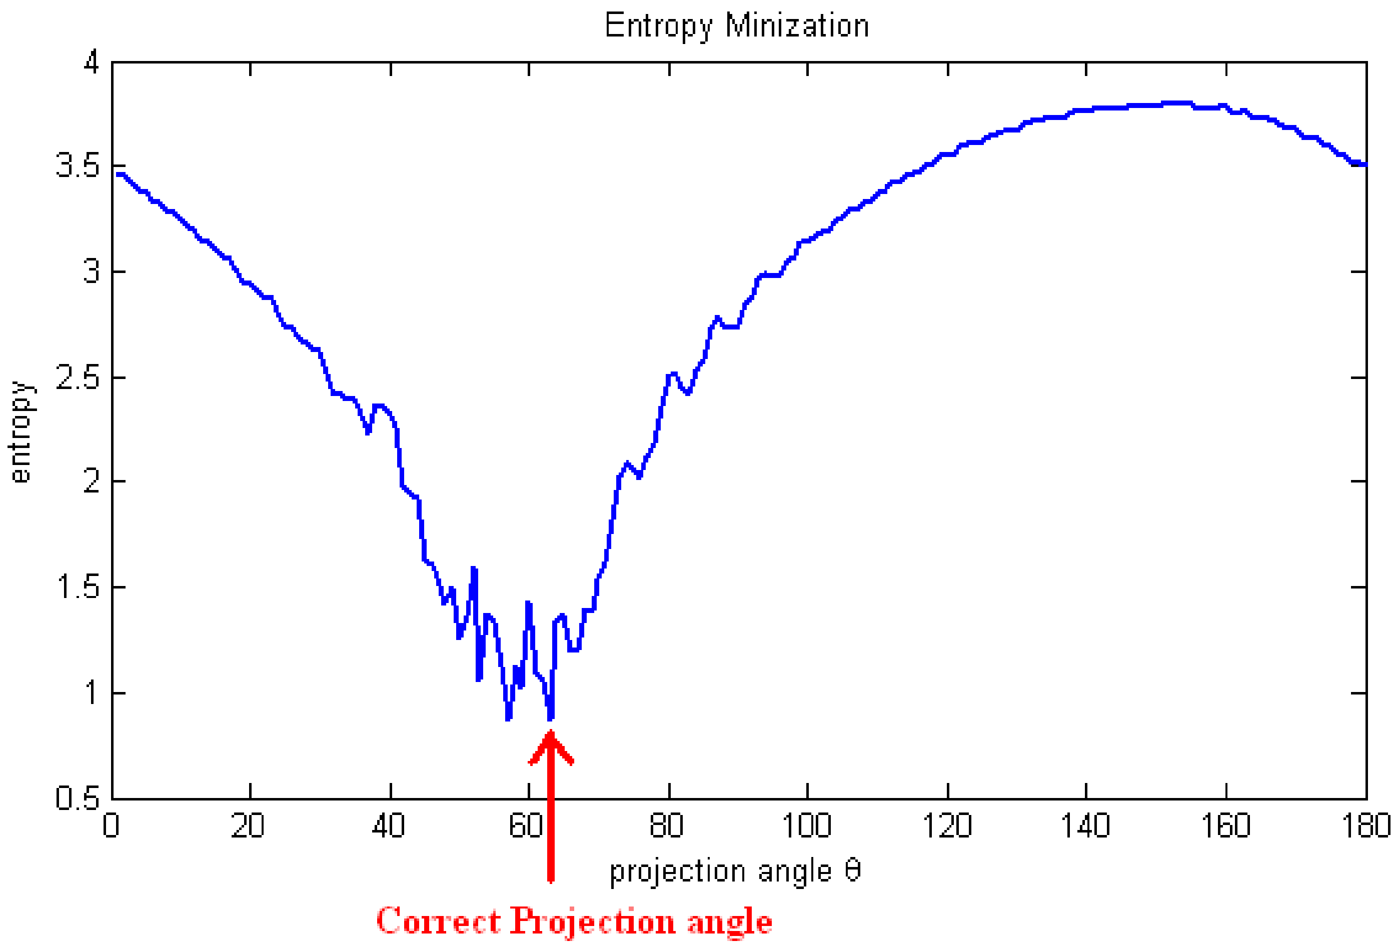

More “spread-out” distribution would produce a larger entropy, hence ***the projection direction that produces the minimum entropy is the correct projection direction***

Sweep Angle of Projection


Recovered Chromaticity

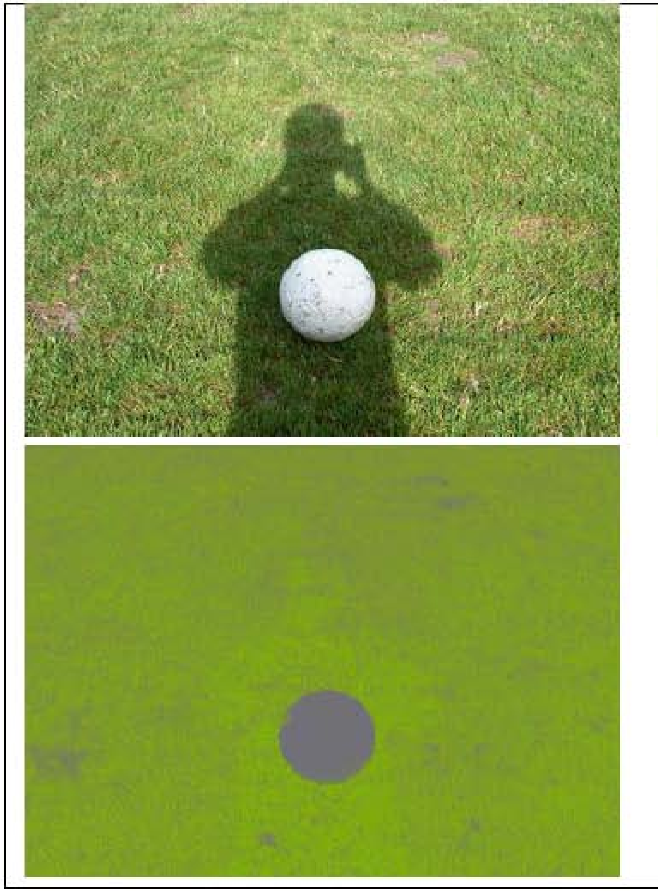

Limitations of Shadow Removal

- Only Hard shadows can be removed
- No overlapping of object and shadow boundaries
- Planckian light sources
- Narrow band cameras are idealized

- Reconstruction methods are texture-dumb

ETH

Discussion…

Assumptions for Both Entropy Min. & Image Seq.

- Each edge can be a shadow border OR a change in the reflectance image
- Remove the shadow edges and get the reflectance image
  - Both algorithms uses the same reconstruction method
- In real life there are soft shadow and vague shadow edges.
  - In many cases there will be mixed edges so the separation can not be easily done
- The second algorithm is more sensitive because it is has less redundancy (one image only!)

Modern Intrinsic Images Refs

- Recovering Reflectance and Illumination in a World of Painted Polyhedra
  - Sinha & Adelson ICCV’93
- Deriving Intrinsic Images from Image Sequences
  - Weiss ICCV’01
- Recovering Intrinsic Images from a Single Image
  - Tappen, Freeman, Adelson NIPS’03
- Intrinsic Images by Entropy Minimization
  - Finlayson, Drew, Lu, ECCV’04

ETH

Property of illuminator: Planck's Law

Planck's Law defines the energy emission rate of a blackbody illuminator , in unit of *watts per square meter per wavelength interval*, as a function of wavelength λ (in meters) and temperature T (in degrees Kelvin).

$$P_r(\lambda) = c_1 \lambda^{-5} \left(e^{\frac{c_2}{\lambda T}} - 1 \right)^{-1}$$

Where $c_1 = 3.74183 \times 10^{-16} \mathrm{Wm}^2$ and $c_1 = 1.4388 \times 10^{-2} \mathrm{mK}$ are constants.

Spectral power:

$$E(\lambda) = I \times P_r = I c_1 \lambda^{-5} \left(e^{\frac{c_2}{\lambda T}} - 1 \right)^{-1} \cong I c_1 \lambda^{-5} e^{-\frac{c_2}{\lambda T}};$$

I is the illumination intensity

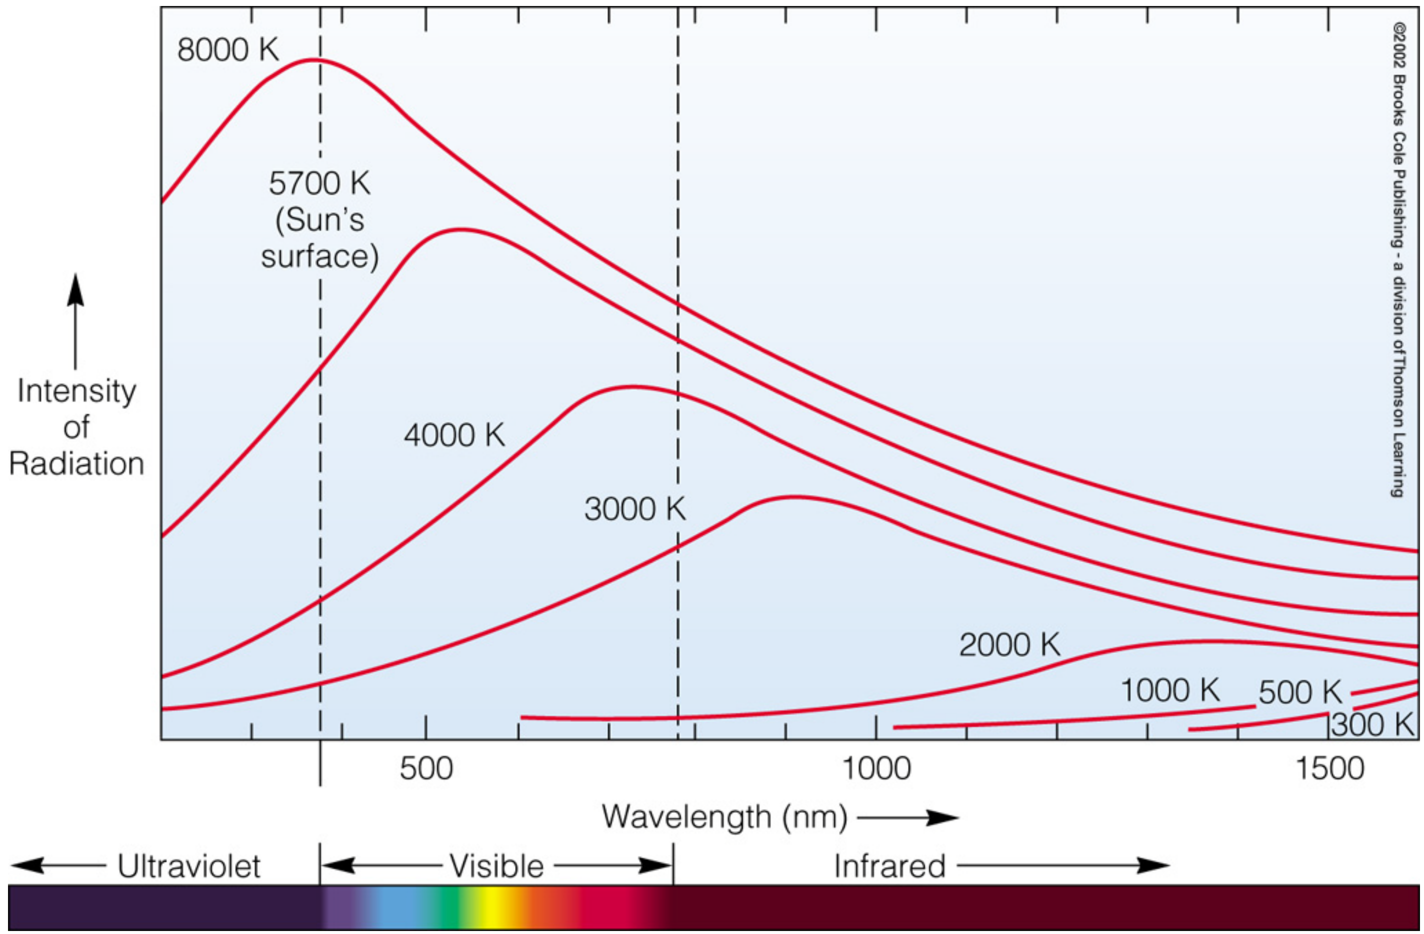

8000 K
5700 K (Sun's surface)
4000 K
3000 K
2000 K
1000 K
500 K
300 K
Intensity of Radiation
500
1000
1500
Wavelength (nm)
Ultraviolet
Visible
Infrared
©2002 Brooks Cole Publishing - a division of Thomson Learning

Calibration results

Nikon CoolPix 995

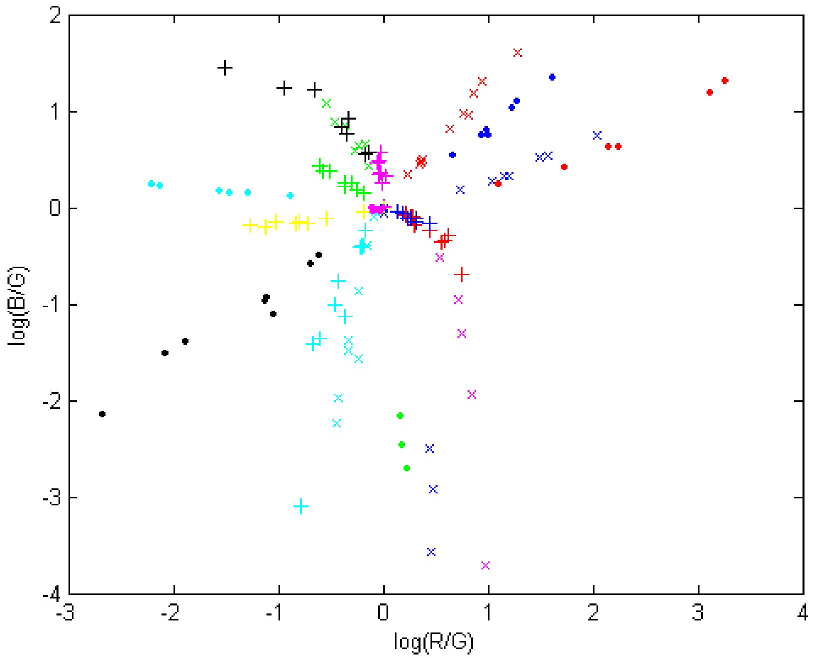

Nikon D-100