Shading and Recognition

OR

The first Mrs Rochester

D.A. Forsyth, UIUC

Structure

- Argument:
  - why shading
  - why shading analysis died
  - reasons for hope
- Some History
  - Primitives
- Reconstructions are possible
  - Variable source shading analysis

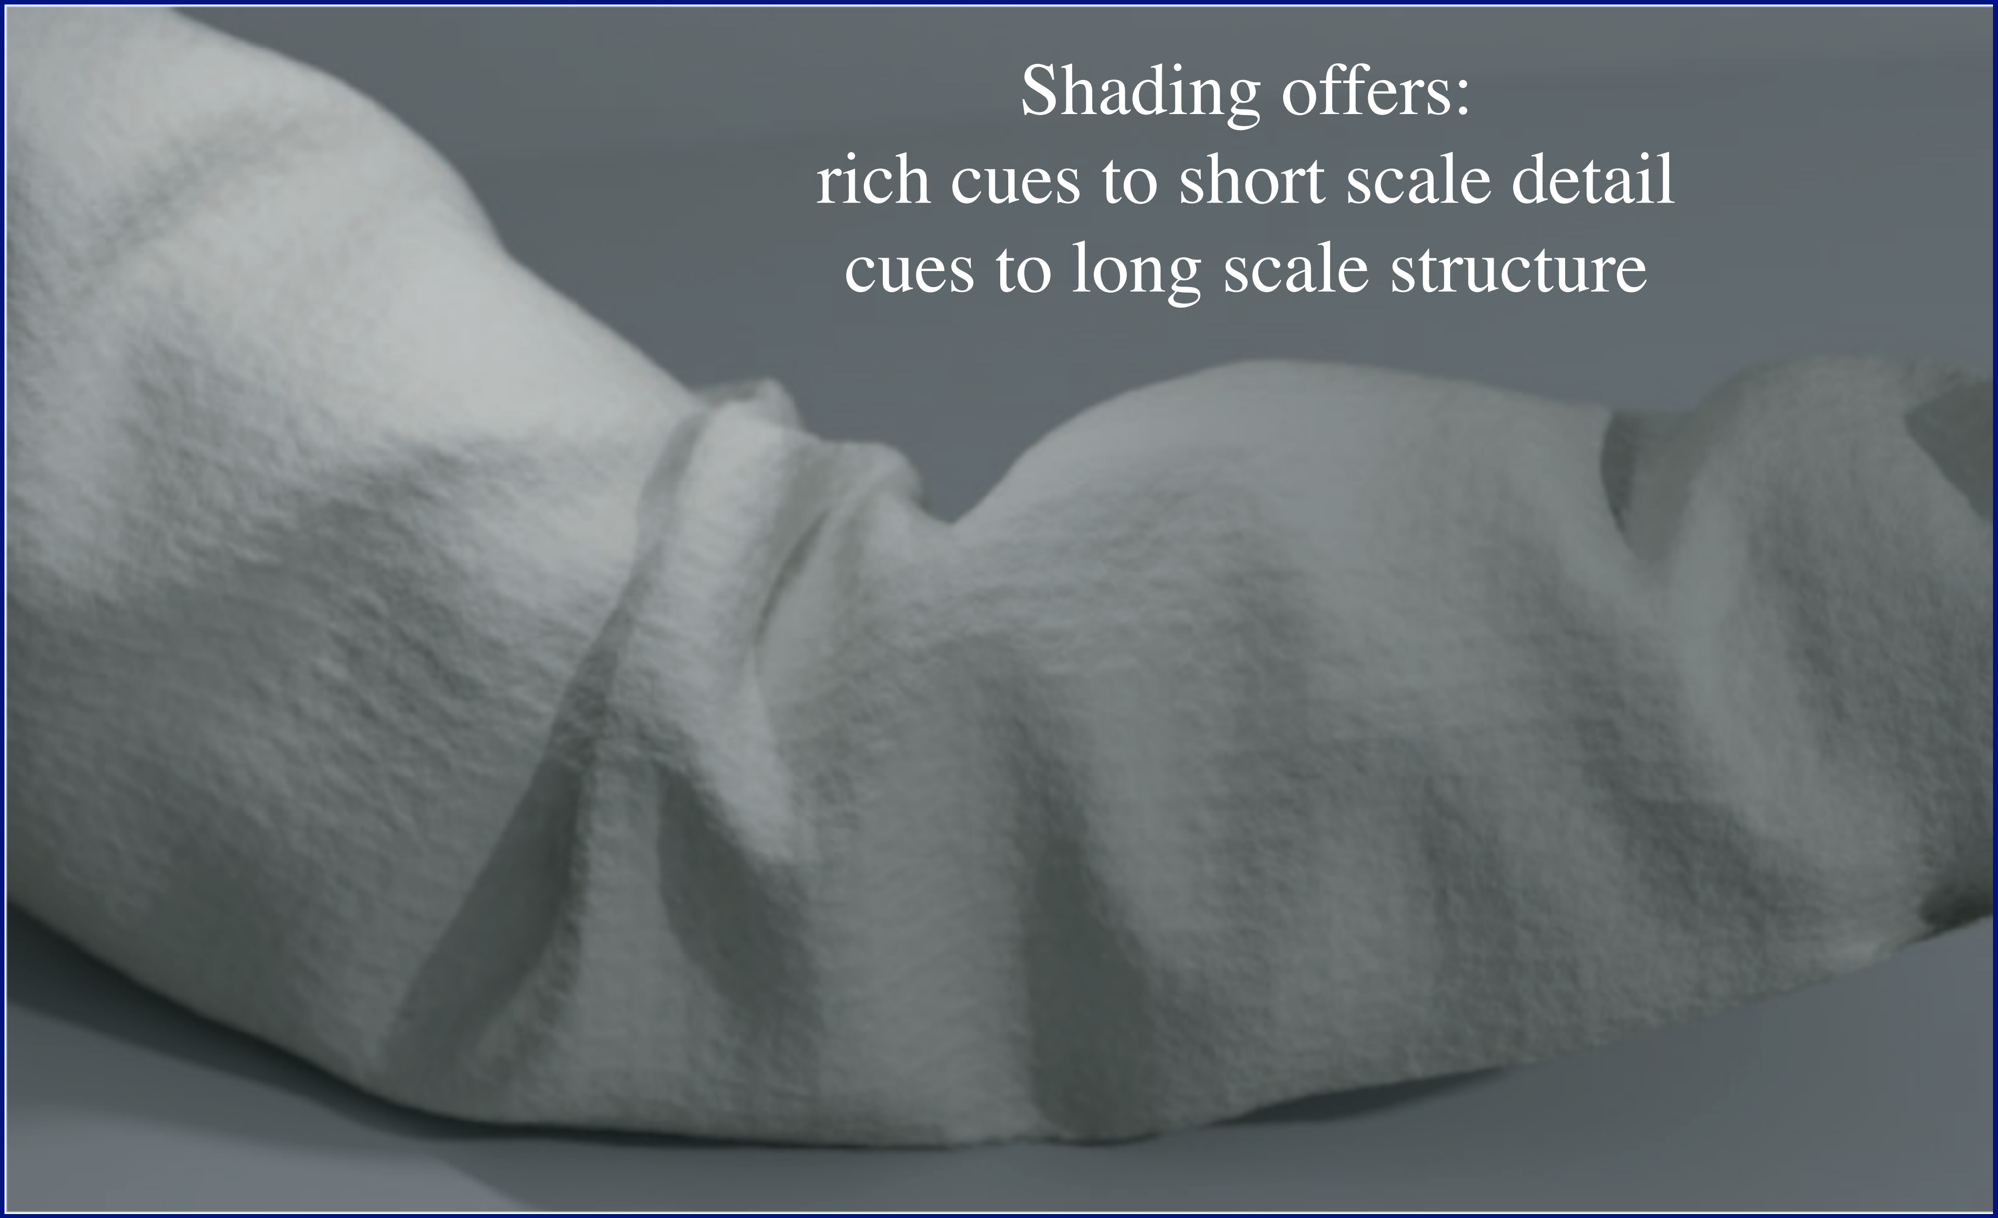

From White+Forsyth 07

Reconstruction from shading

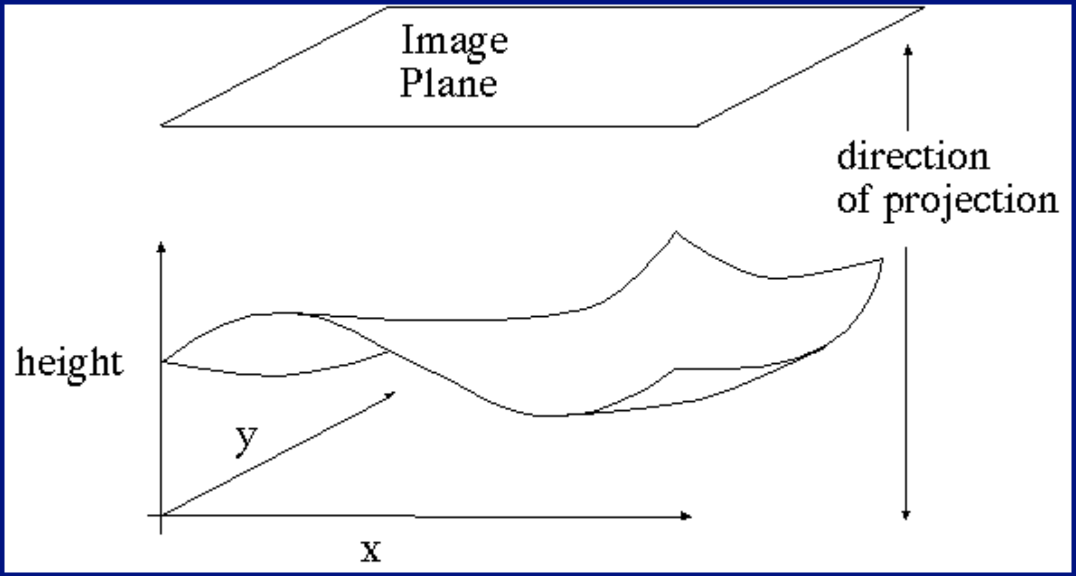

- Conventions:
  - Orthography
    - (but. for example, Prados+Faugeras
  - Height field
    - partial derivatives are written p, q

Reconstruction from shading

$$R(p, q; \mathbf{S}) = I(x, y)$$

Reflectance Map | Image intensity

- Local model
  - Points with the same normal get the same shading value
- The Image Irradiance Equation (IIE)
  - Horn, 1970 and lots of later papers by lots of authors
- This is a PDE
  - First order, non-linear, actually Hamilton Jacobi

Physical Critiques

- Real shading is not local
  - interreflections
    - points with the same normal get different shading values
- Devastating
  - because a physically exact formulation is unmanageable
    - (it has been tried, Nayar et al 91)
    - cannot account for distant radiators we can't see

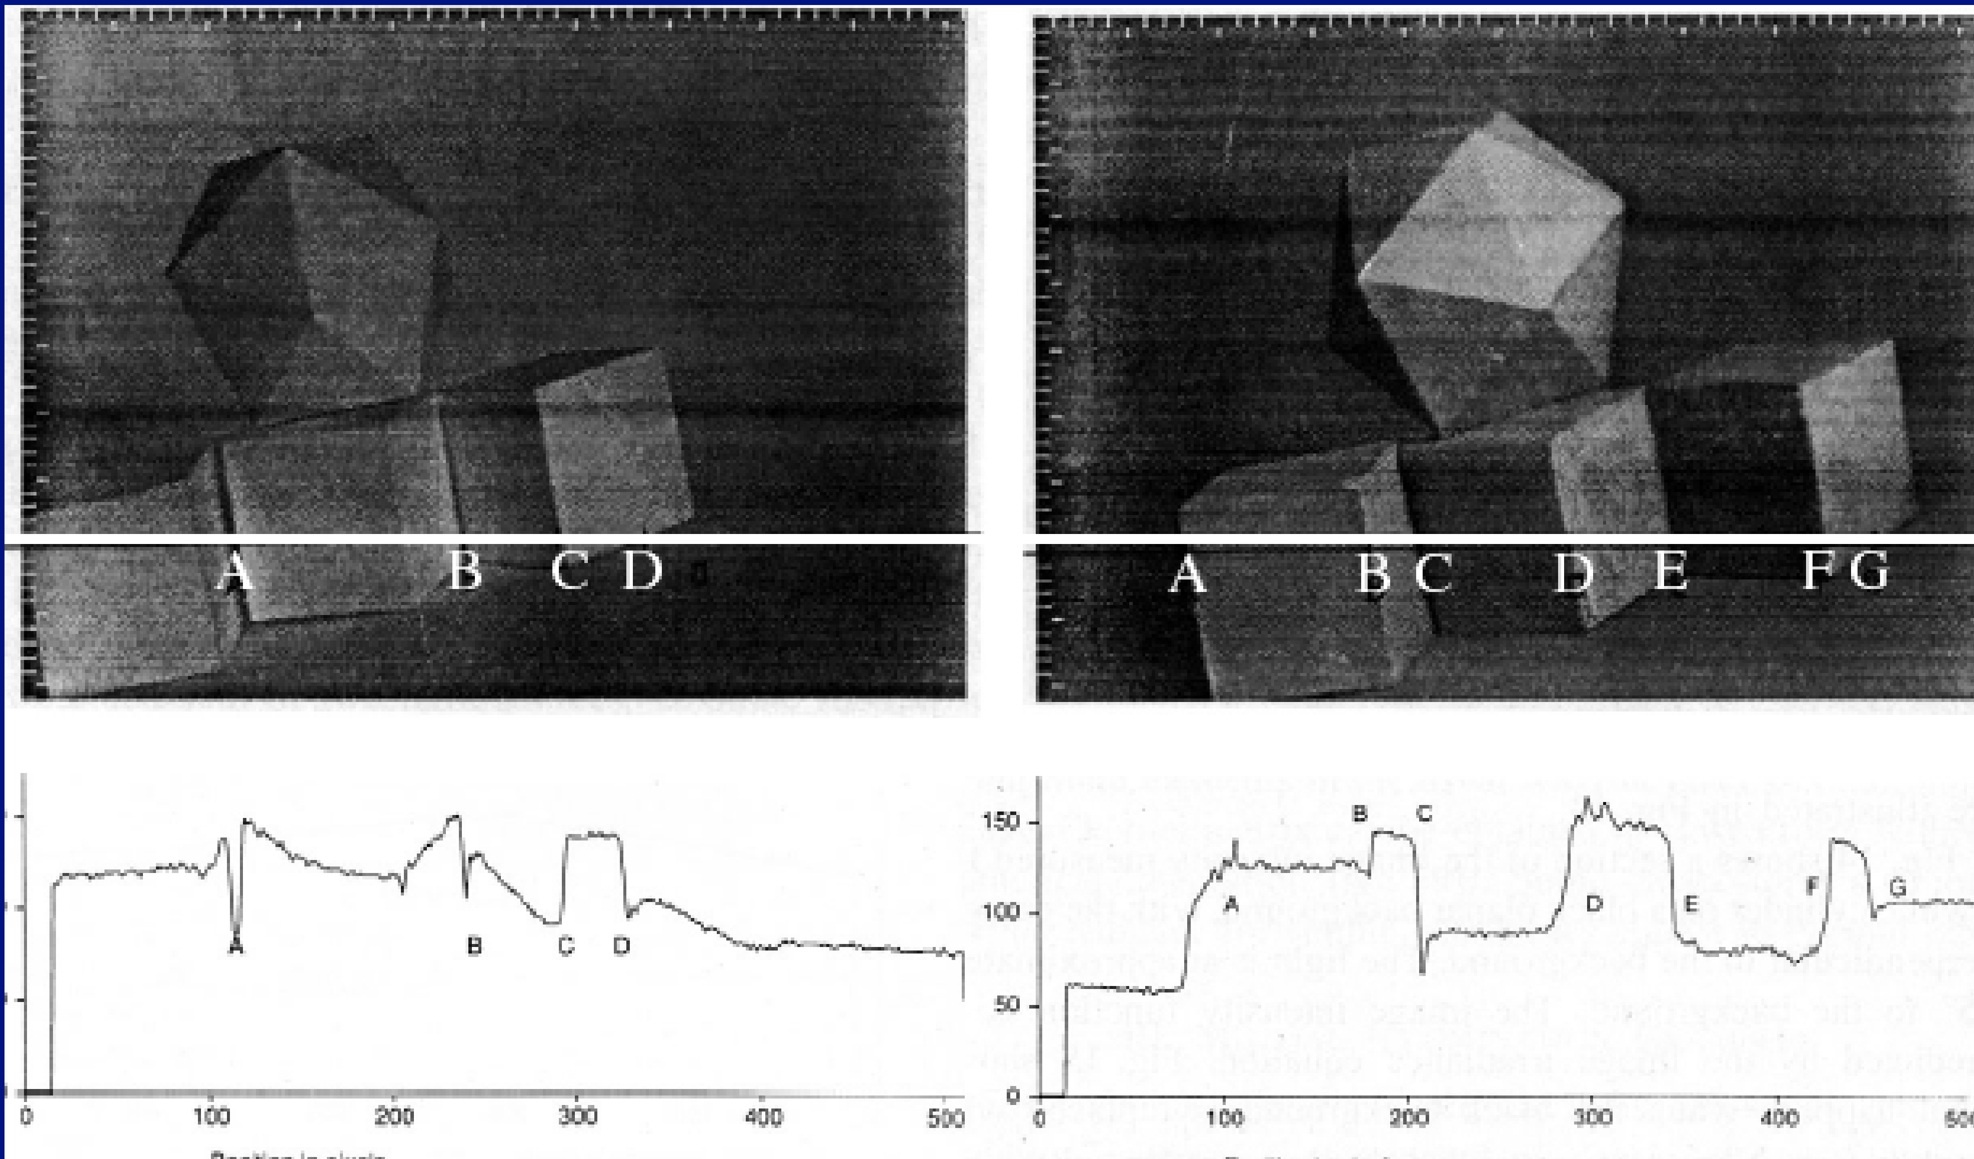

Forsyth Zisserman ’89, ’91
after Gilchrist, Koenderink, etc.

Existence

- For some boundary conditions, a solution to IIE exists
  - but for relatively few;
    - as geometric constraints grow stronger, existence almost always fails
    - by easy characteristic strip argument

- Characteristic strip
  - FO PDE is ODE, along CS
  - CS given by ODE
  - CS's don't intersect

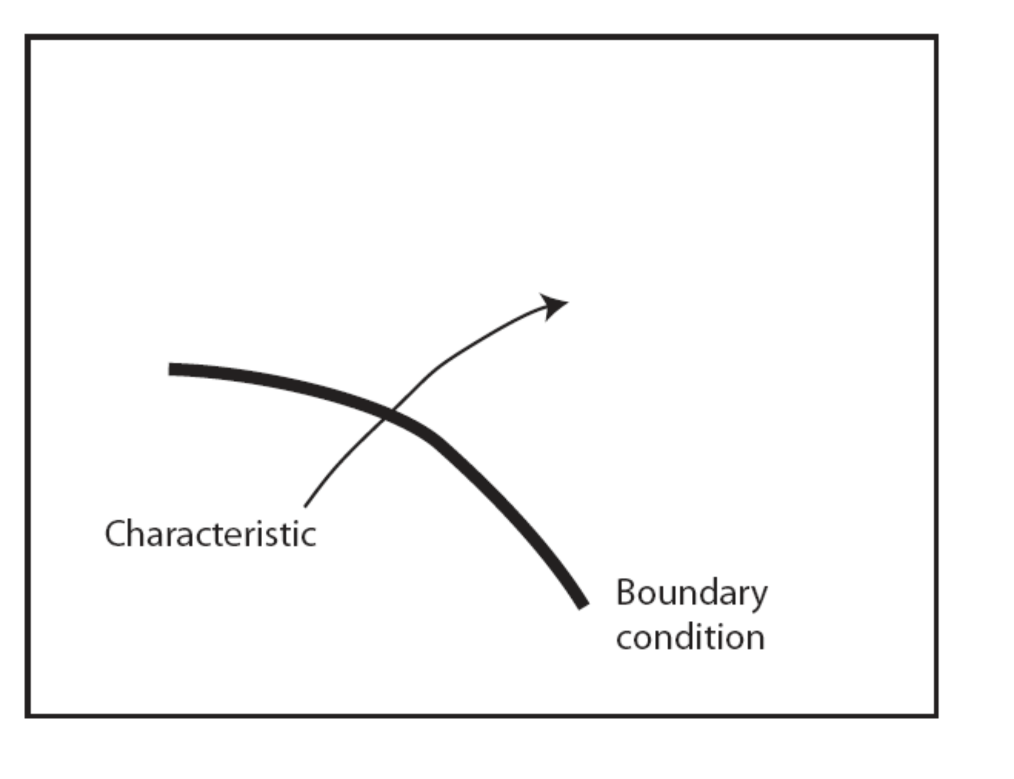

More BC’s -> Existence fails

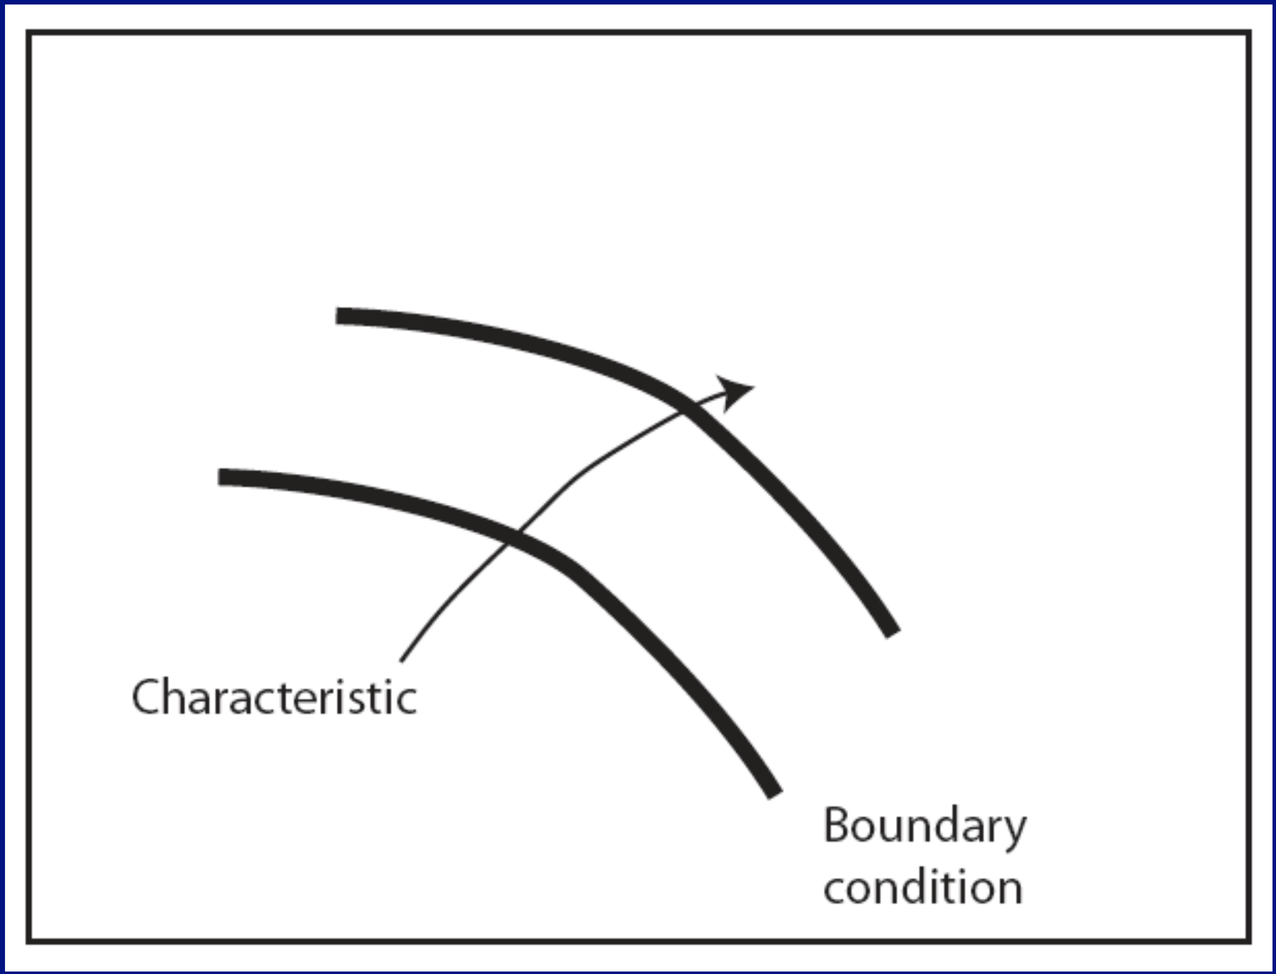

Existence

- Solutions do not exist for rich boundary conditions
  - current literature says:
    - not a problem - want reconstruction from minimal geometry data
- Options
  - classical fails
  - Lipschitz (too many solutions)
  - Viscosity (one, but no physical justification for choice)
    - RouyTourin 92, Lions et al 93, Prados Faugeras 03
- Real world
  - many rich sources of geometric constraint (identity; stereo; SFM;...)
  - should not impede existence

General Guideline: A formulation which doesn't have existence for natural problem instances needs to be fixed

Relations to photometric stereo

- Photometric stereo
  - estimate normal (+albedo) locally from multiple shading maps
  - using completely different procedures, reasoning

Pragmatics

- Shape from shading doesn't work
  - ample evidence
    - No comparison between right answer and reconstructions
    - Poor results on synthetic (!) data

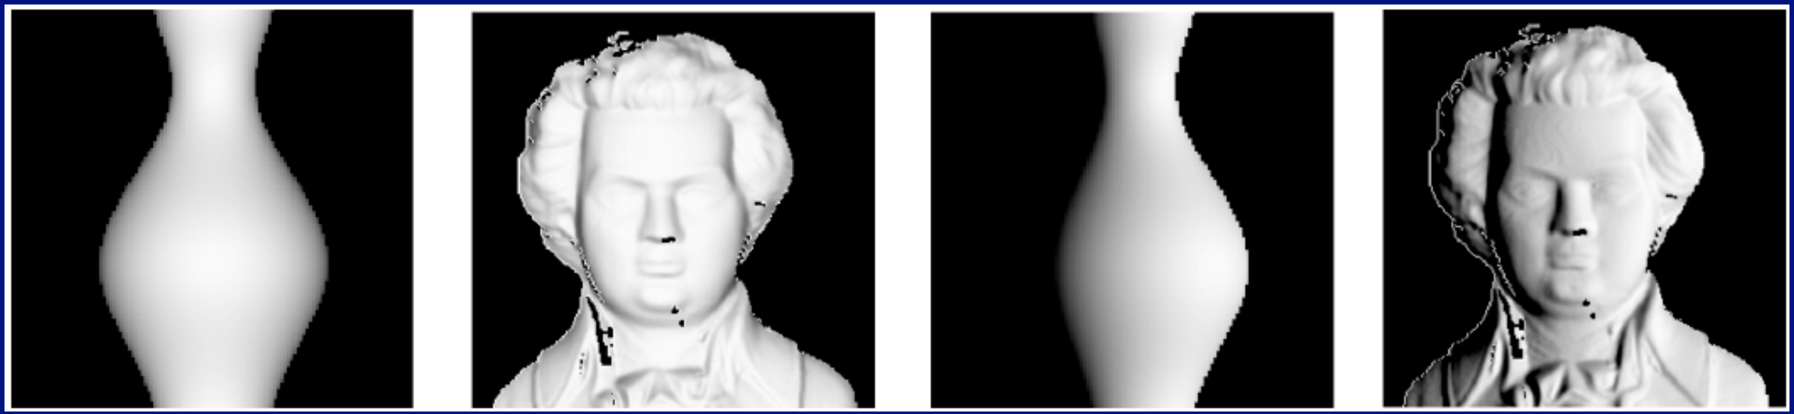

From Zhang ea, 99

Pragmatics

From Zhang ea, 99

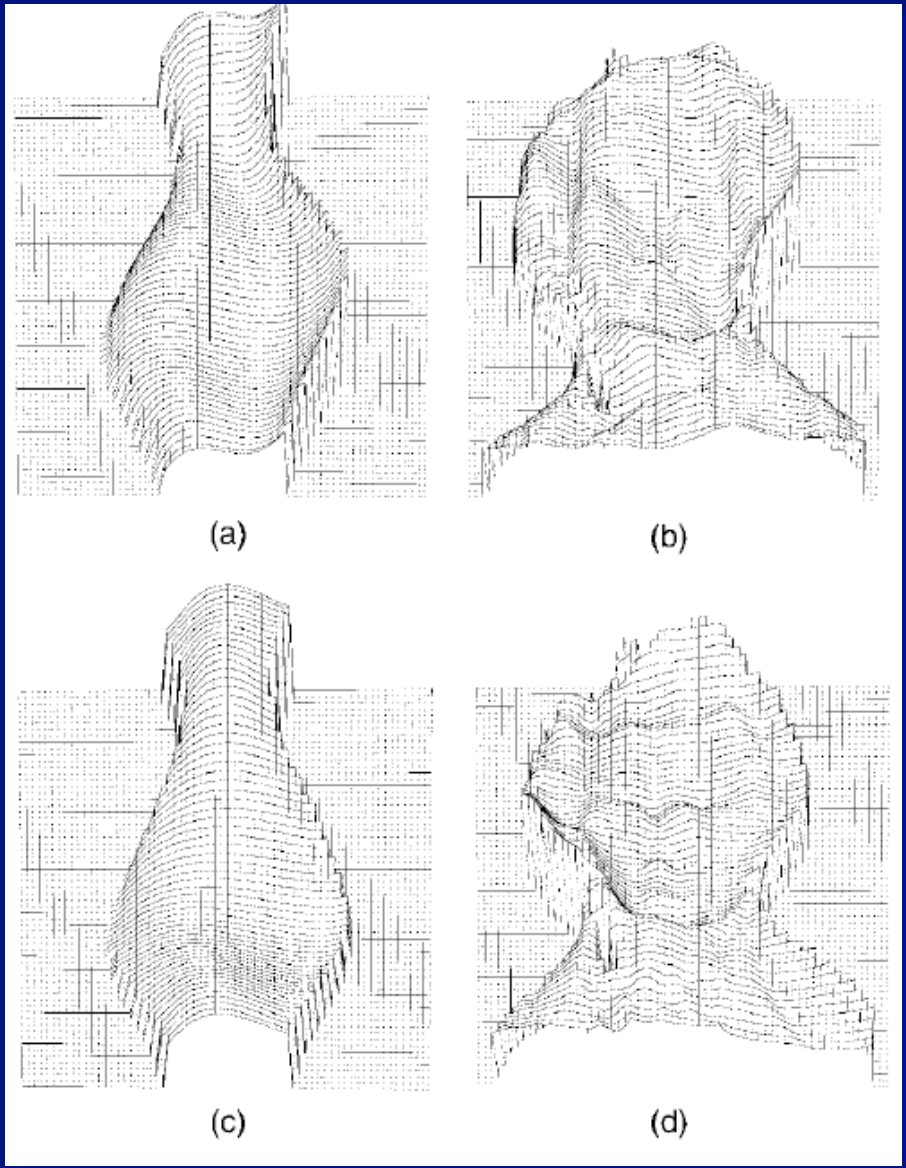

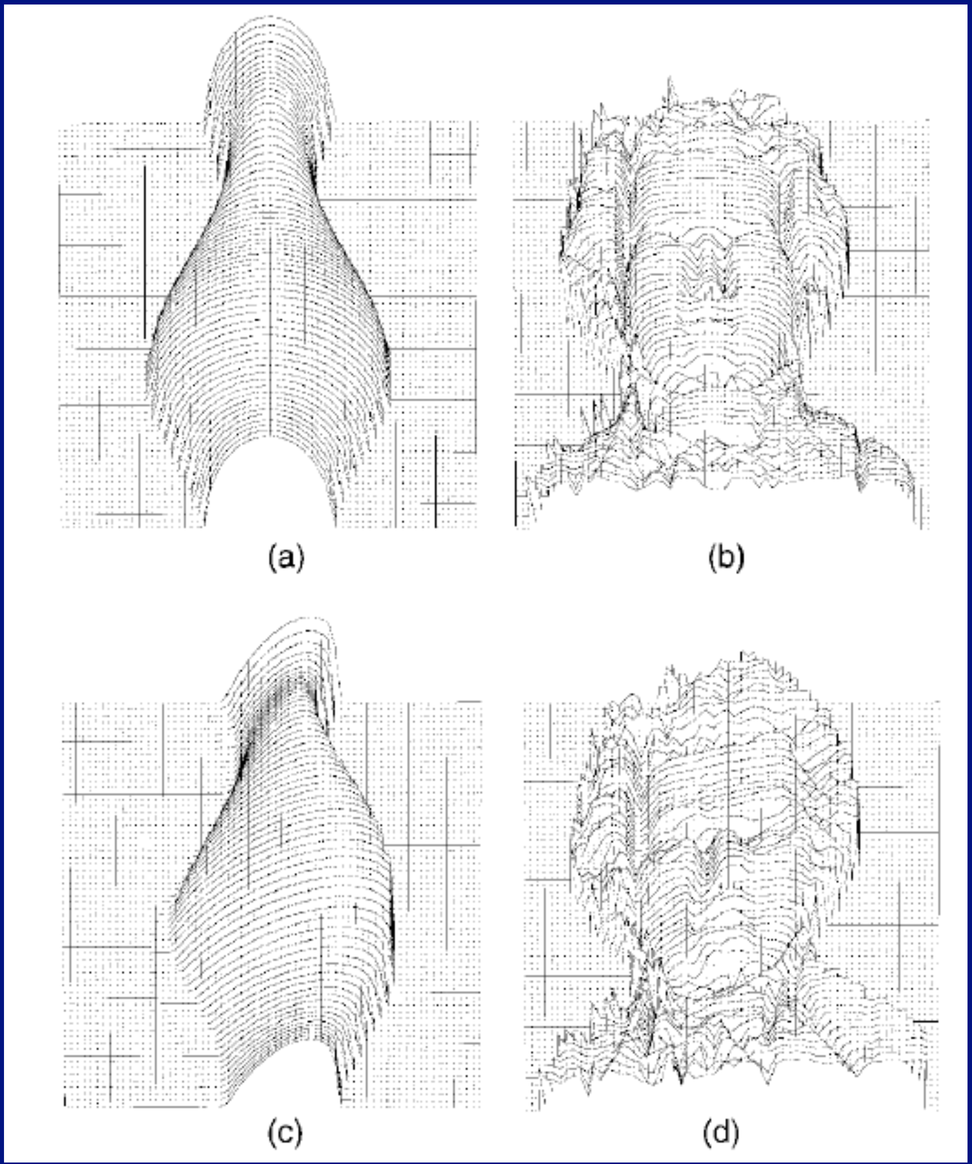

Pragmatics

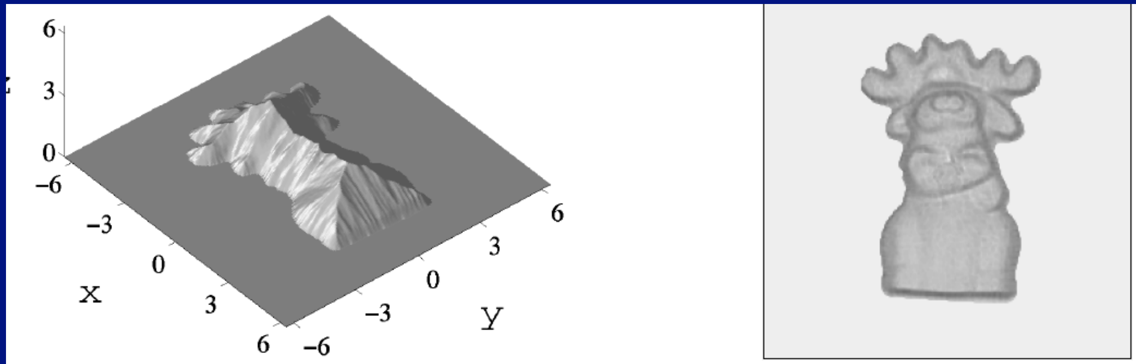

Fig. 26. FS on Elk: 0.42 *s*.

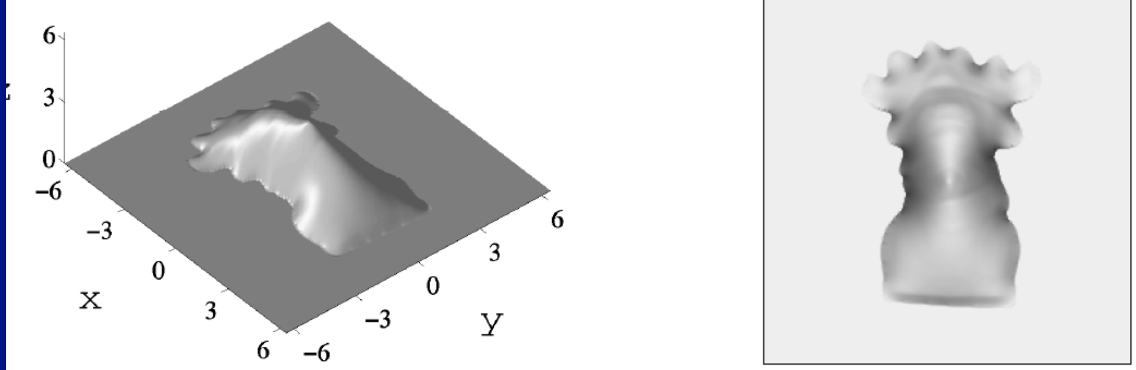

Fig. 27. DD on Elk: 1.25 *s*.

Durou et al, 2007

Pragmatics

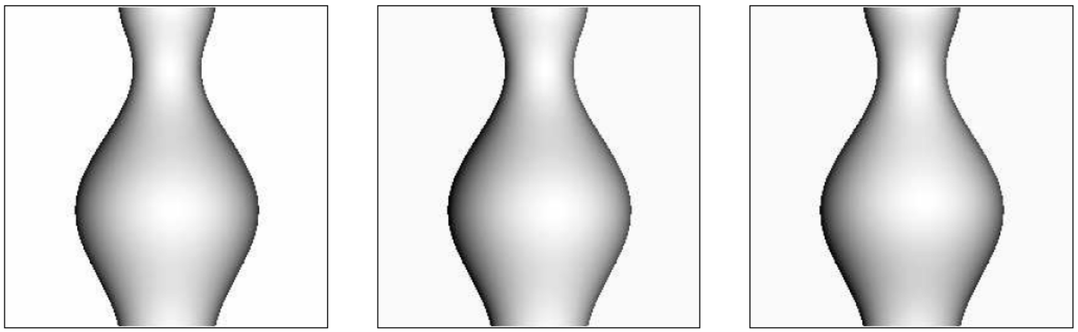

(a) (b) (c)

Figure 31: SV 256×256 images: (a) $\omega = \omega^1$, (b) $\omega = \omega^2$ and (c) $\omega = \omega^3$.

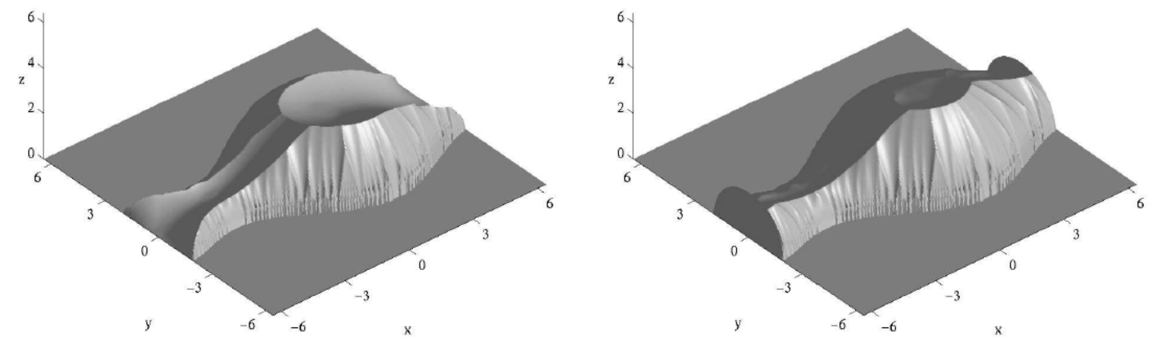

(a) (b)

FS on SV ($\omega = \omega^1$): computed shapes with (a) $u = 0$ and (b) $u = g_{\text{SV}}$ on the boundary.

Durou et al, 2007

Minor critiques

- The world isn't ideal diffuse
  - True, but so what - if we can't solve the easiest case...
- There are specularities
  - see above
    - and we can build specularity detectors
- Albedo varies
  - but we have quite good theories of how to infer albedos

Reasons for hope

- Evidence for pragmatic information in shading
  - SF(T+S)
- Evidence that shading cues are compelling to humans
  - Textureshop
  - Retexturing movies
  - Complex, mixed picture from psychophysics
- Evidence that shading is distinctive
  - Face detectors
  - Some others, rather ragged

SF(T+S) Shading disambiguates texture

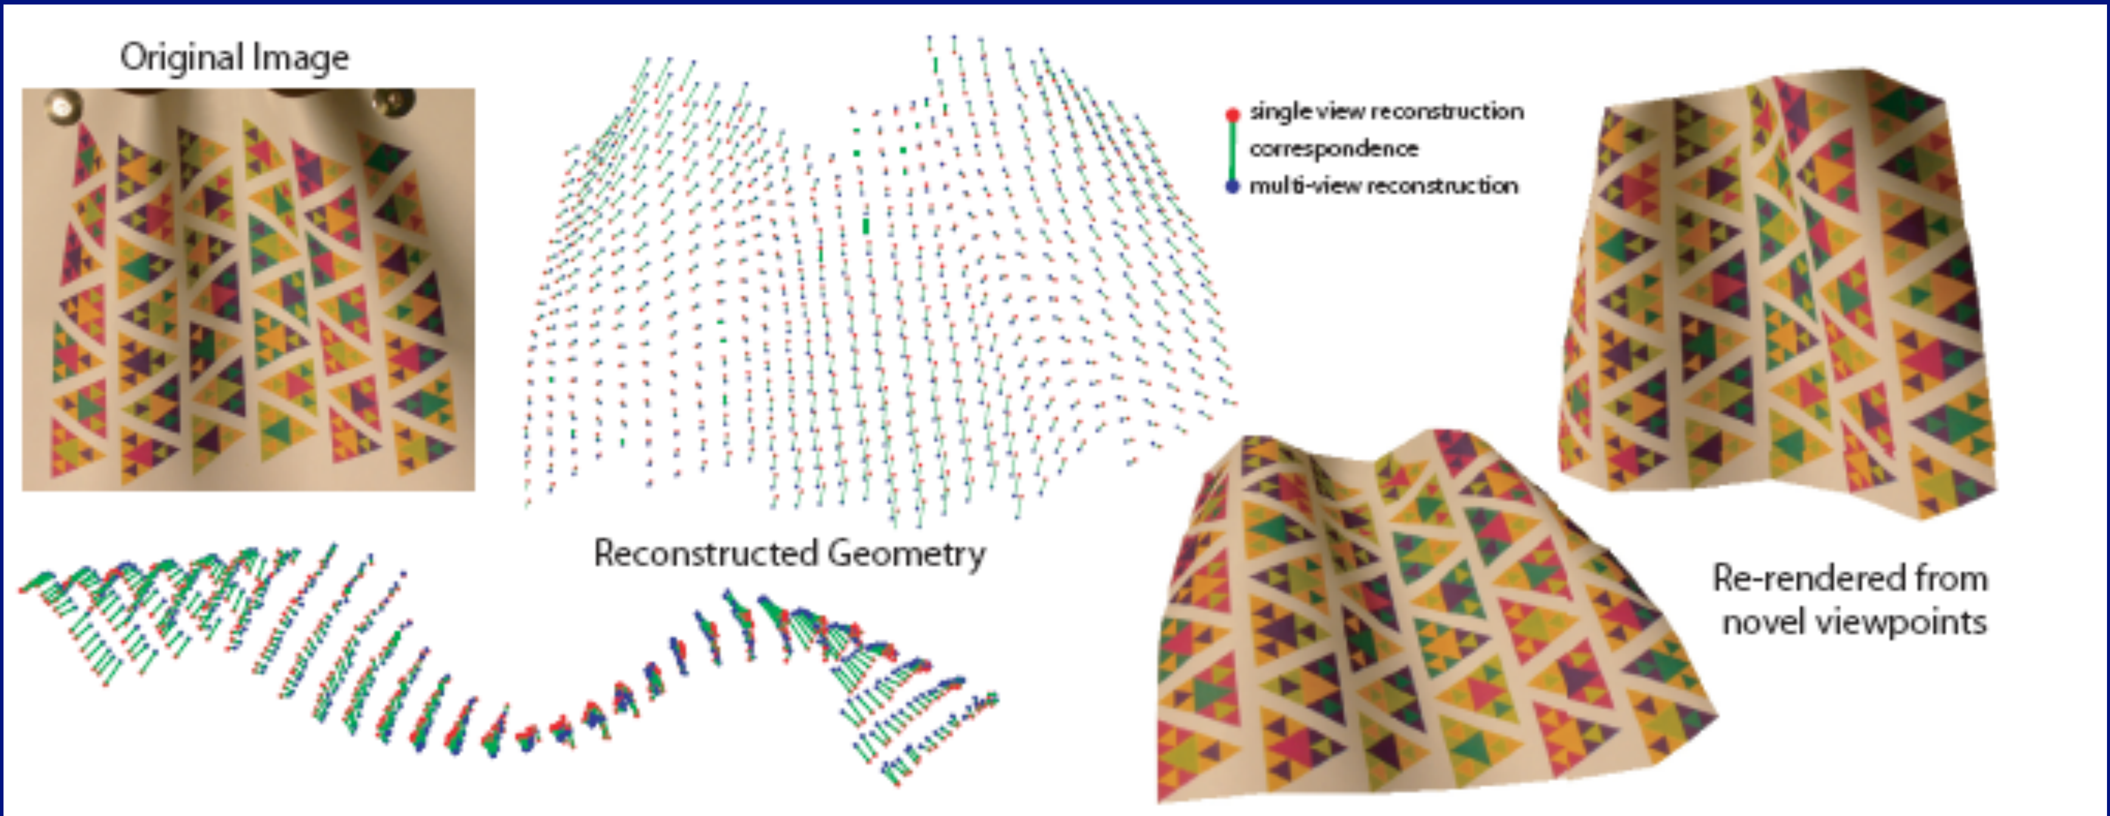

White+Forsyth 06

Textureshop

- Hart+Fang, 04
- Retexture illuminated surface by:
  - Obtaining normal estimate from local shape from shading
    - normal estimate is largely meaningless
  - Use this to compute texture normal
  - Shade this texture with original illumination estimate
- Interesting because
  - In a cue conflict between texture and shading, texture loses

Fang + Hart 04

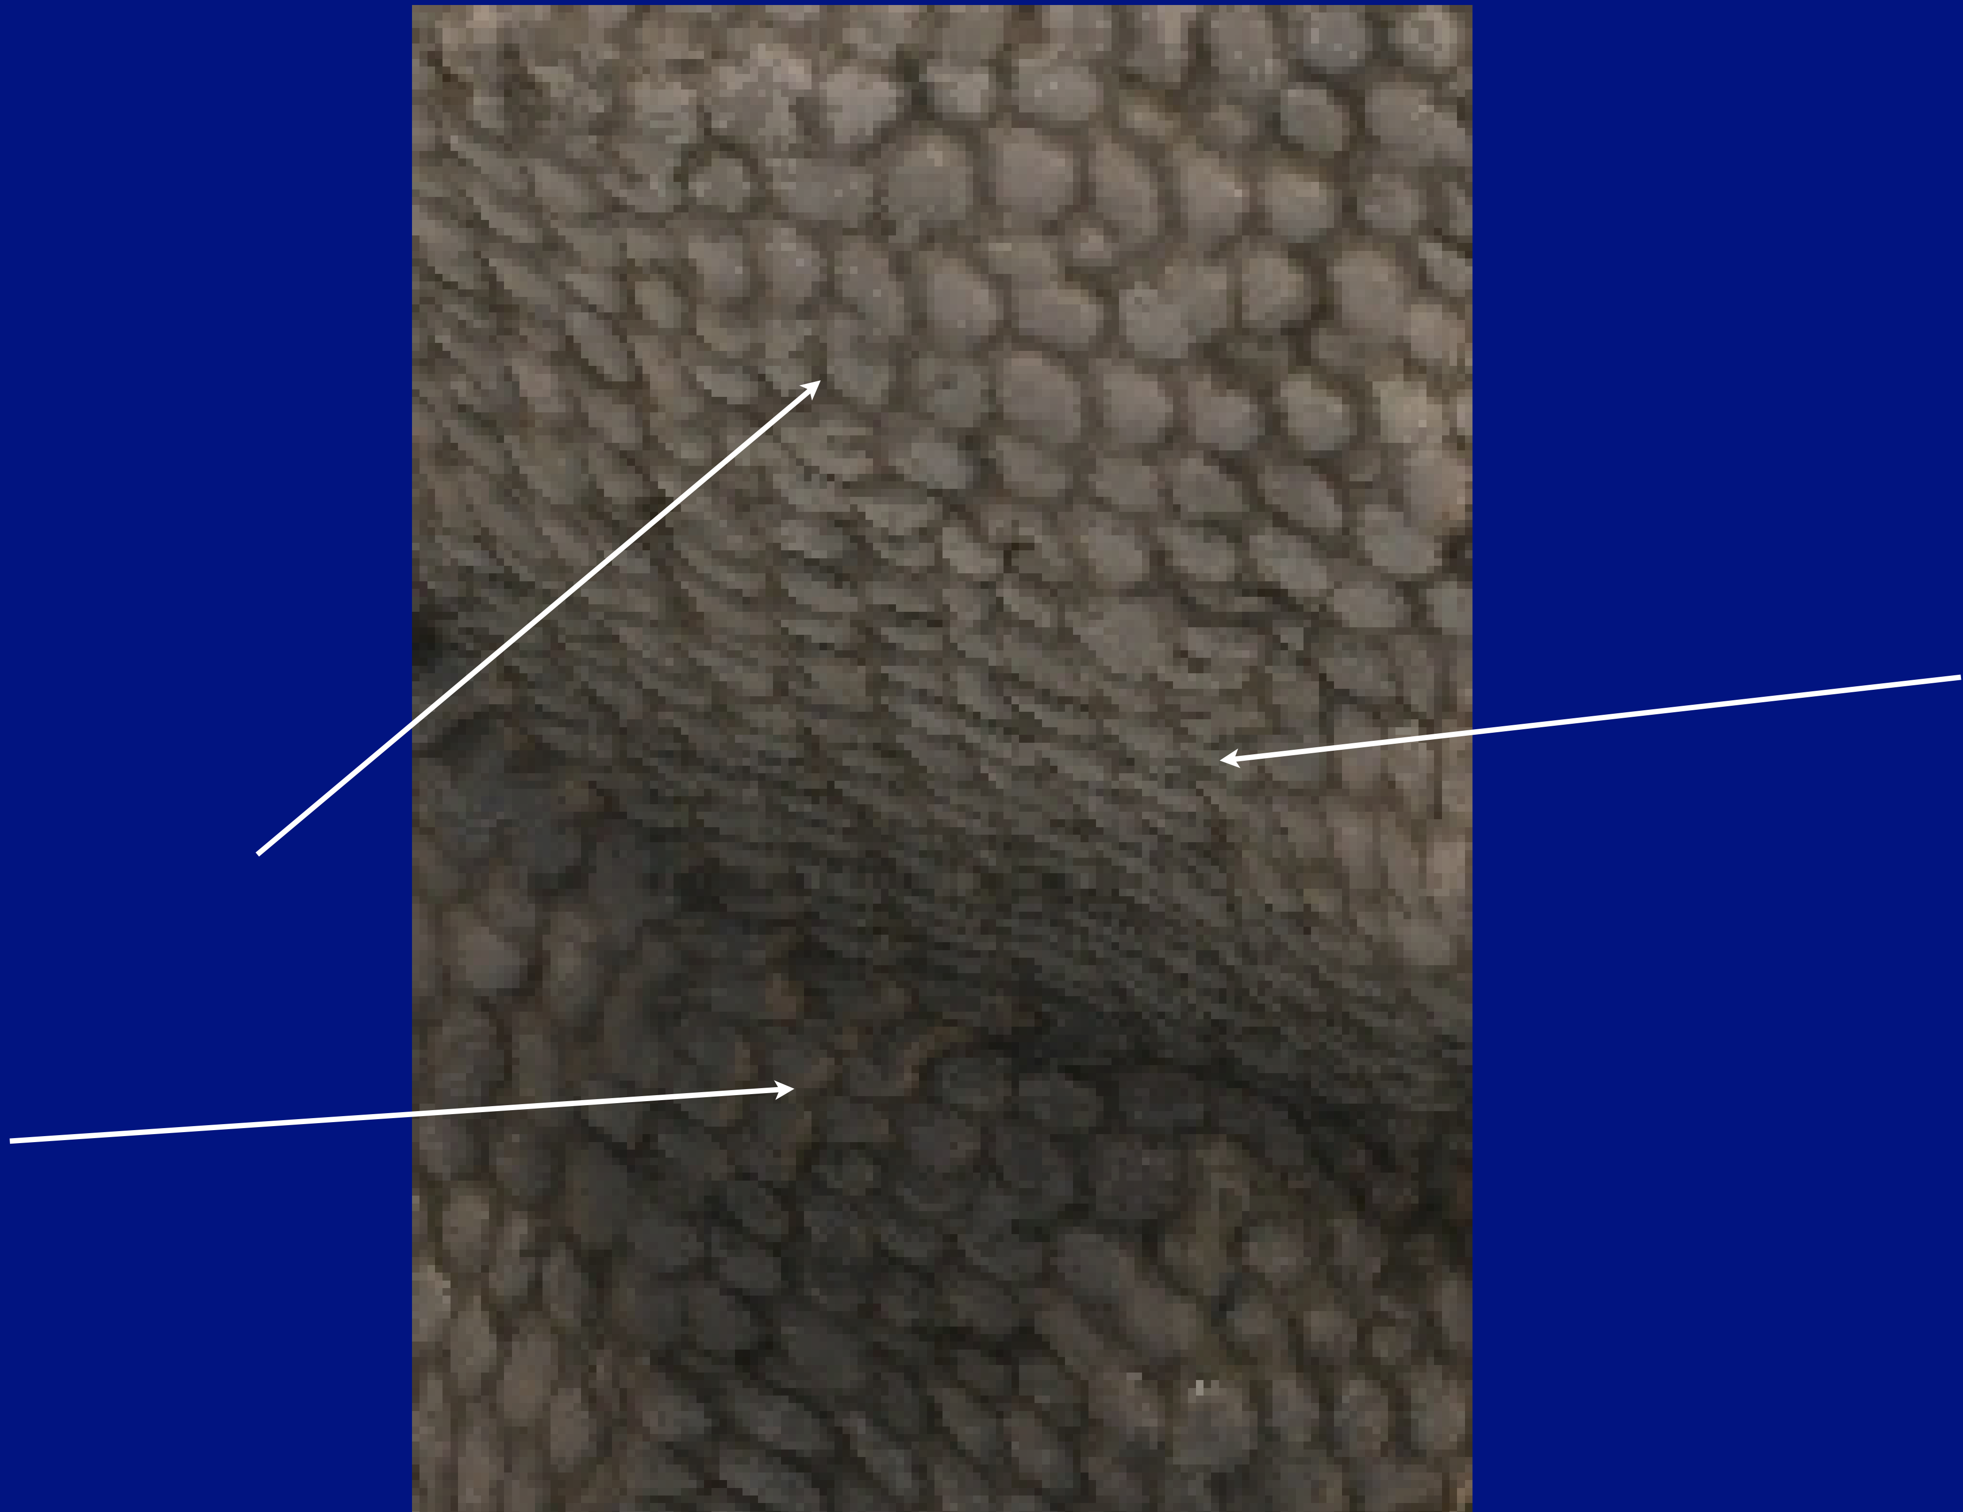

Retexturing movies

- White+Forsyth 06
- Retexture moving surfaces by
  - Building non-parametric estimate of illumination from corners
    - assuming silkscreen, known colors, not known texture
  - Rectify texture to very rough geometric (affine distortion) model
  - Shade with illumination estimate
- Get shading right, it looks natural with weak geometry
  - Shading cues beat motion cues? (at short scales?)
  - Quality issues are
    - flicker
    - surfaces look rigid when fold shading is not reproduced.

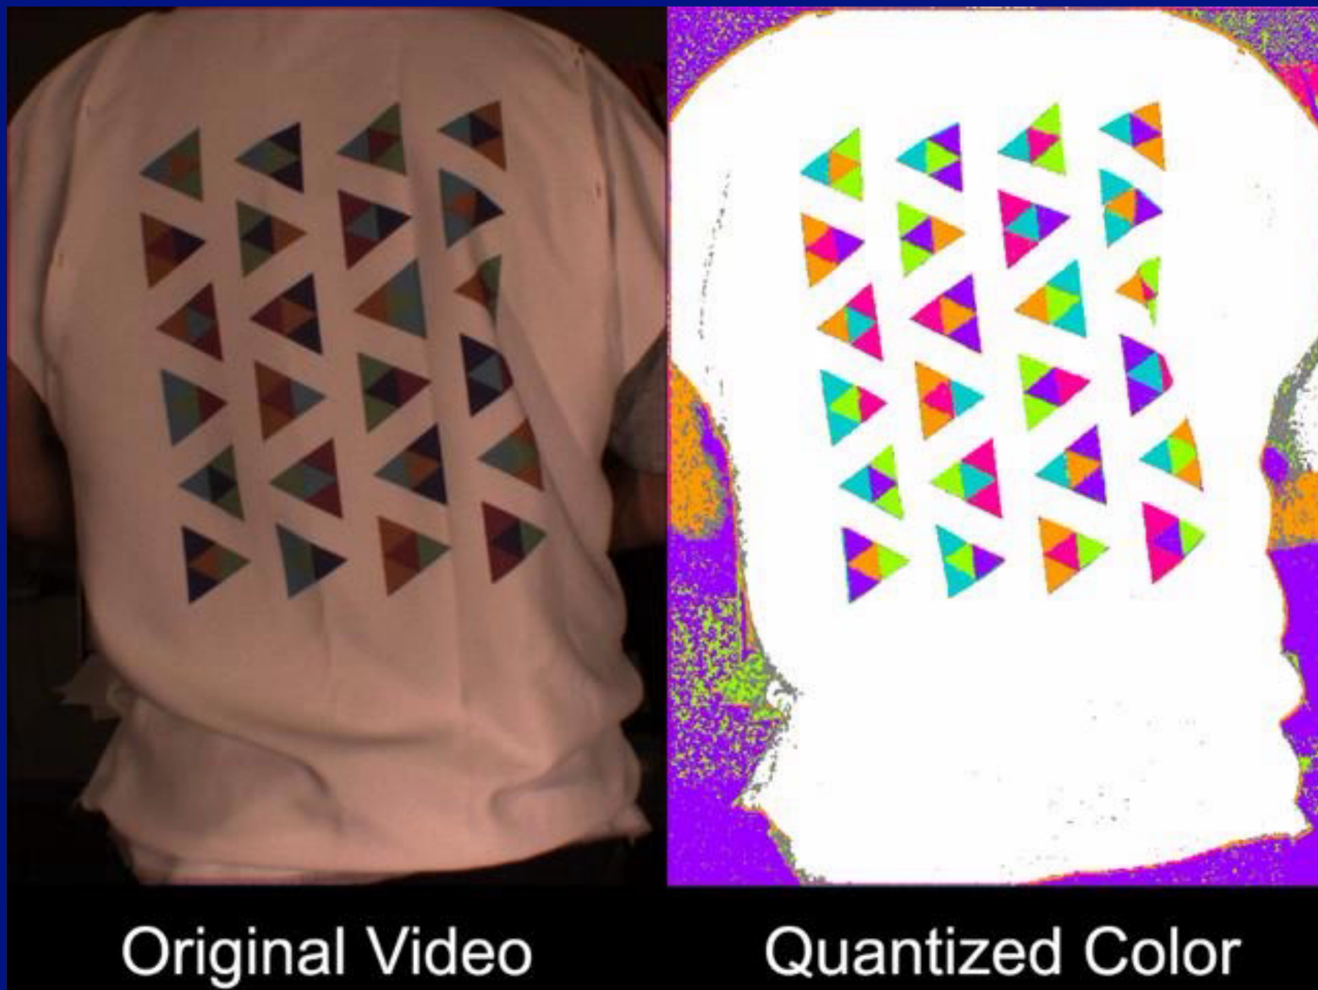
Original Video
Quantized Color

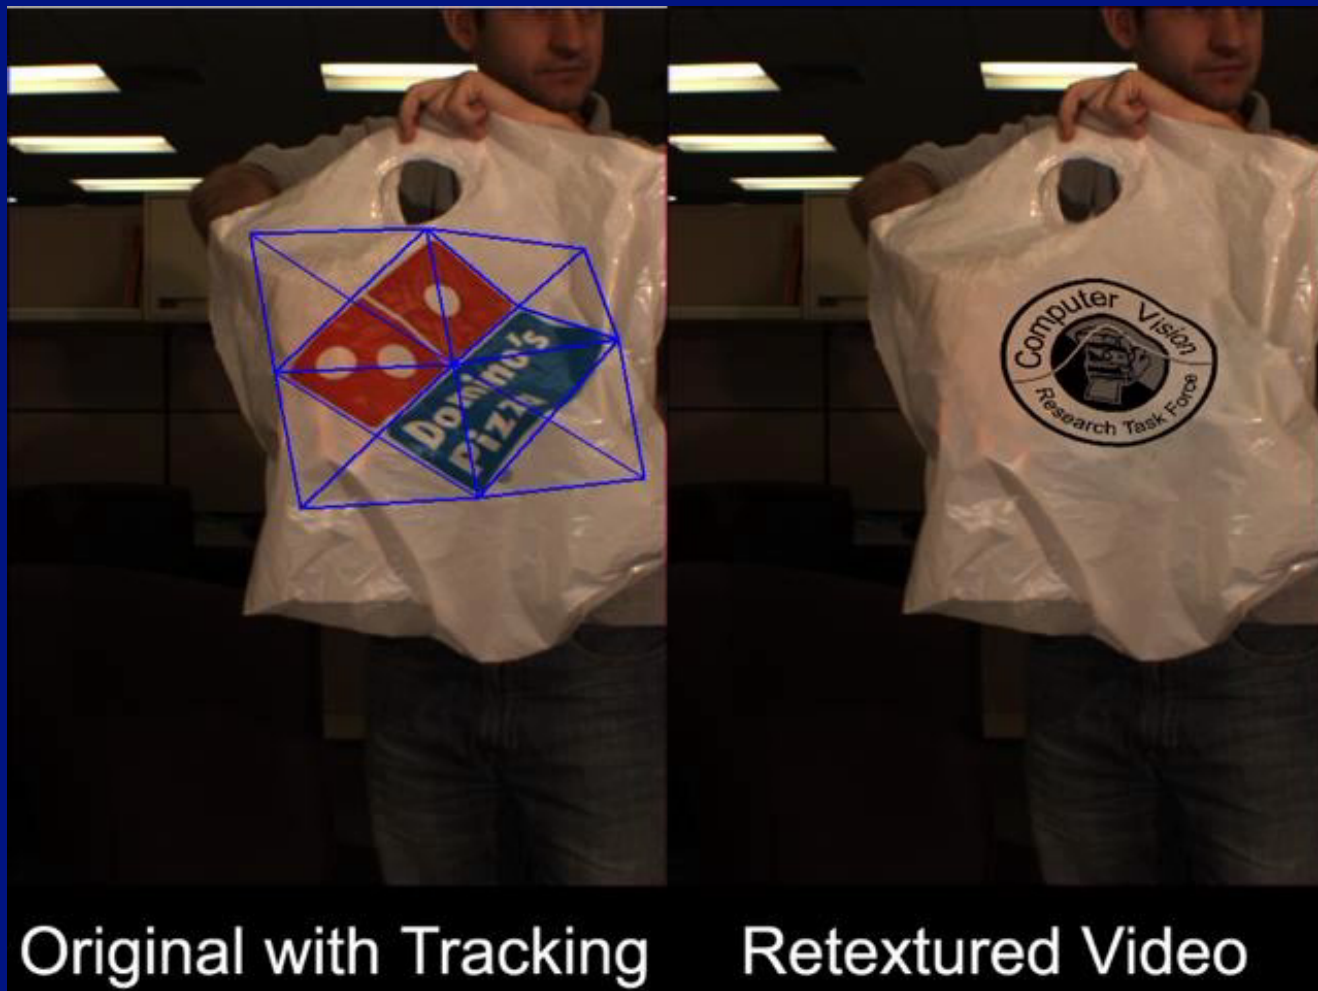
Domino's Pizza
Computer Vision
Research Task Force
Original with Tracking
Retextured Video

Original Video Quantized Color

Structure

- Argument:
  - why shading
  - why shading analysis died
  - reasons for hope
- History
  - Primitives
- Reconstructions are possible
  - Variable source shading analysis

Shading Primitives

- Shading patterns on certain structures are stylized
  - We might be able to spot such patterns and use them
- Huge success
  - Frontal face detectors
- But...
  - few examples
    - Pits, etc. (Koenderink ’83)
    - Folds, Grooves, Cylinders (HaddonForsyth, 98a, b)
    - Objects in fixed configuration (Belhumeur+Kriegman ’98)
  - hard to deploy in natural ways

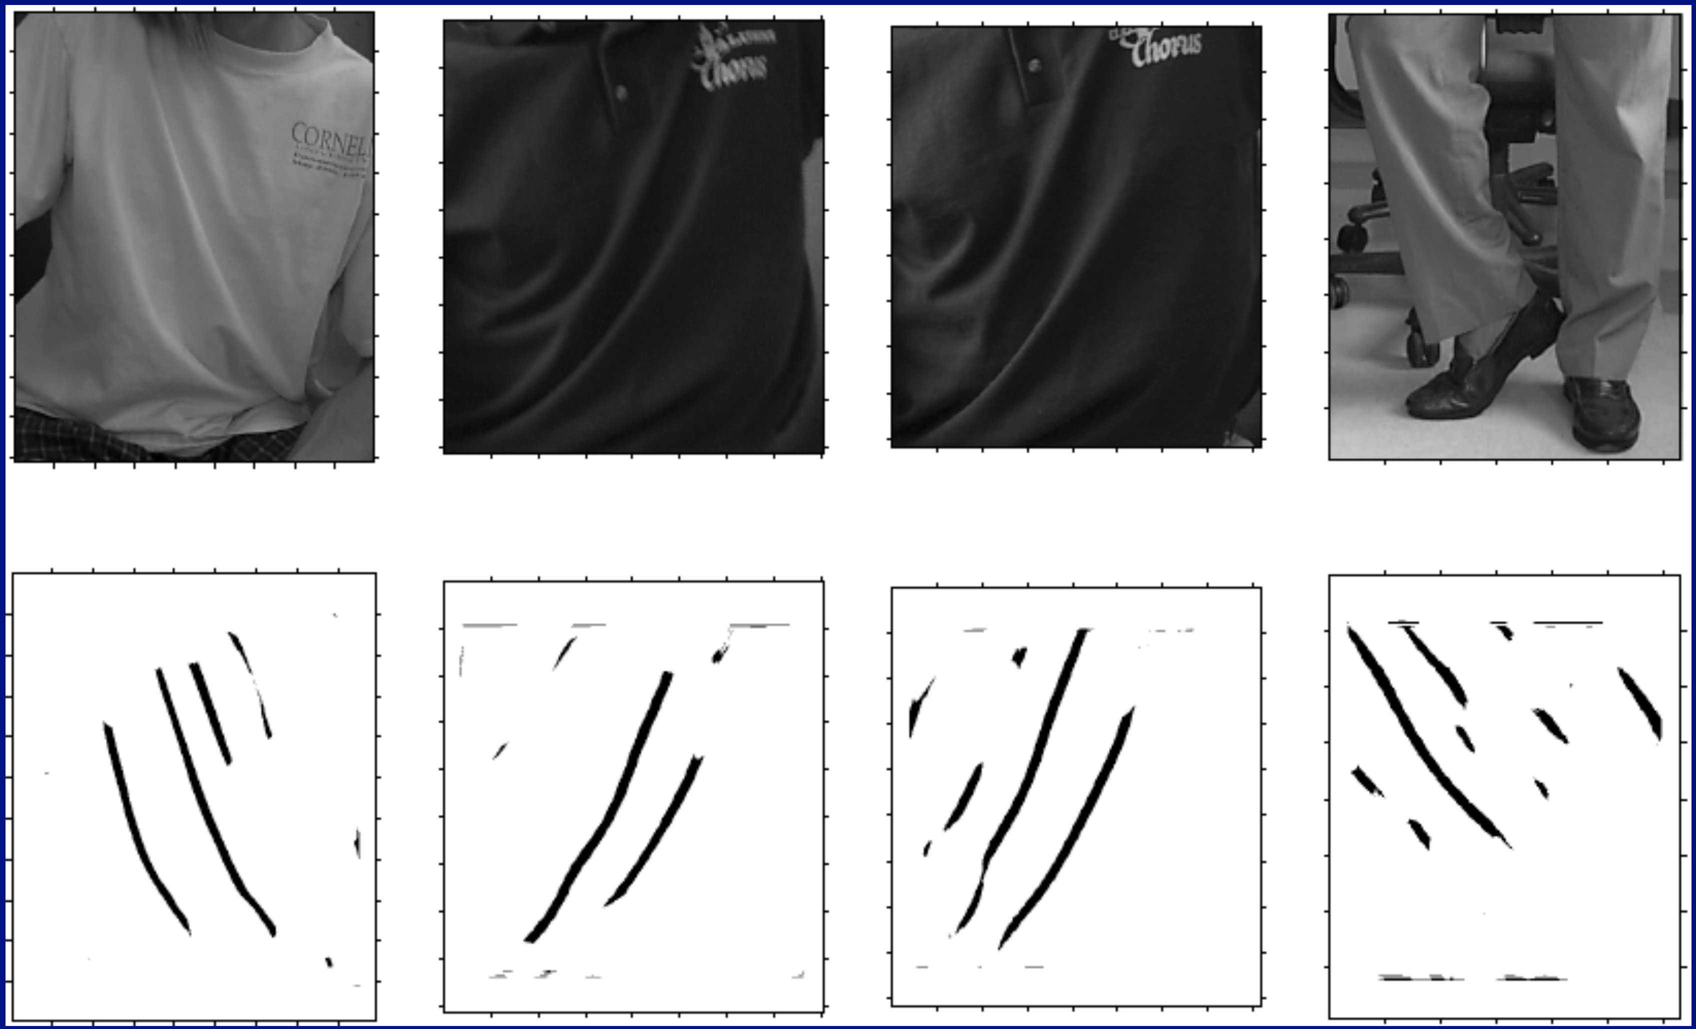

HaddonForsyth 98a

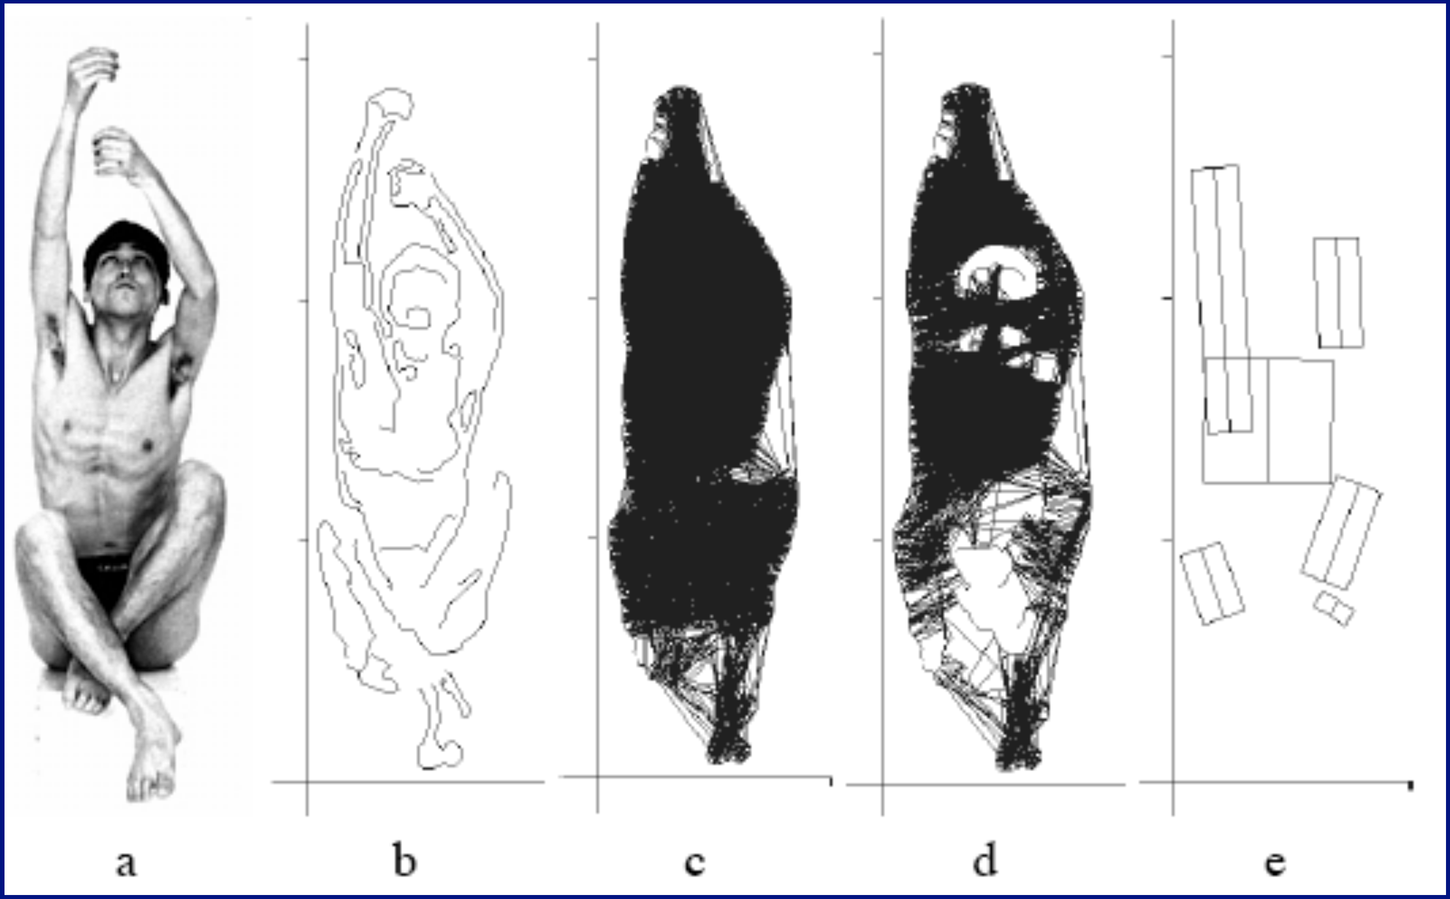

HaddonForsyth 98b

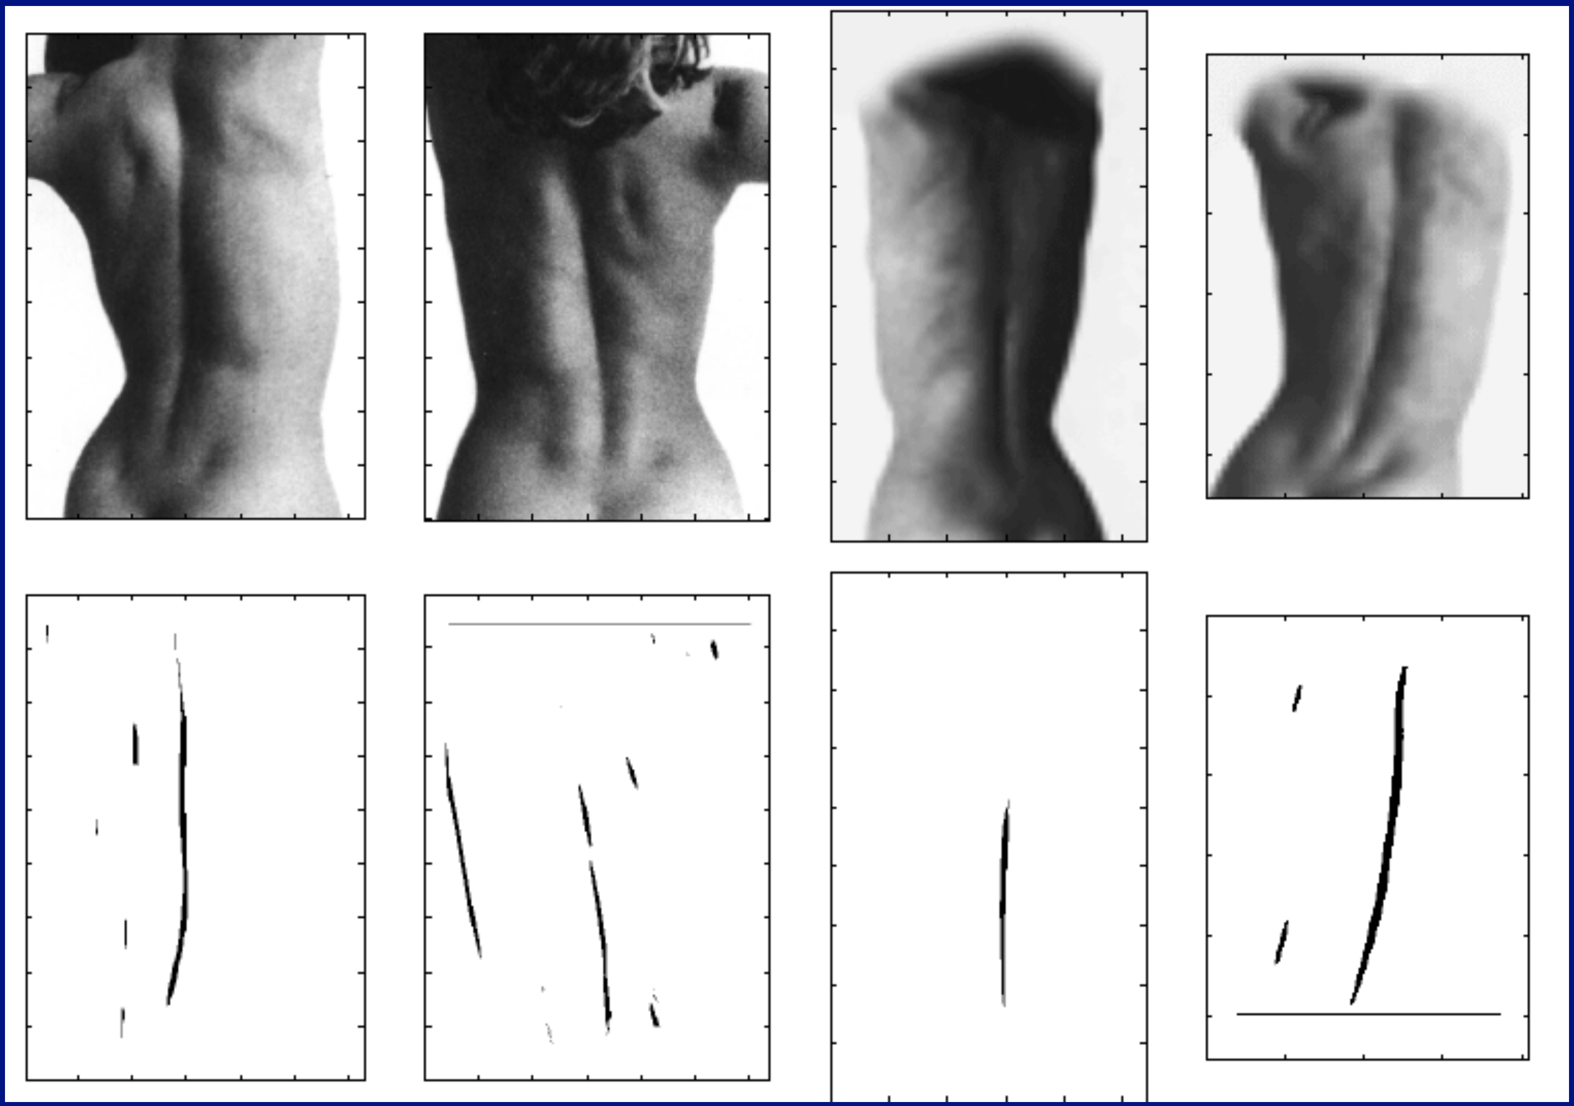

HaddonForsyth 98b

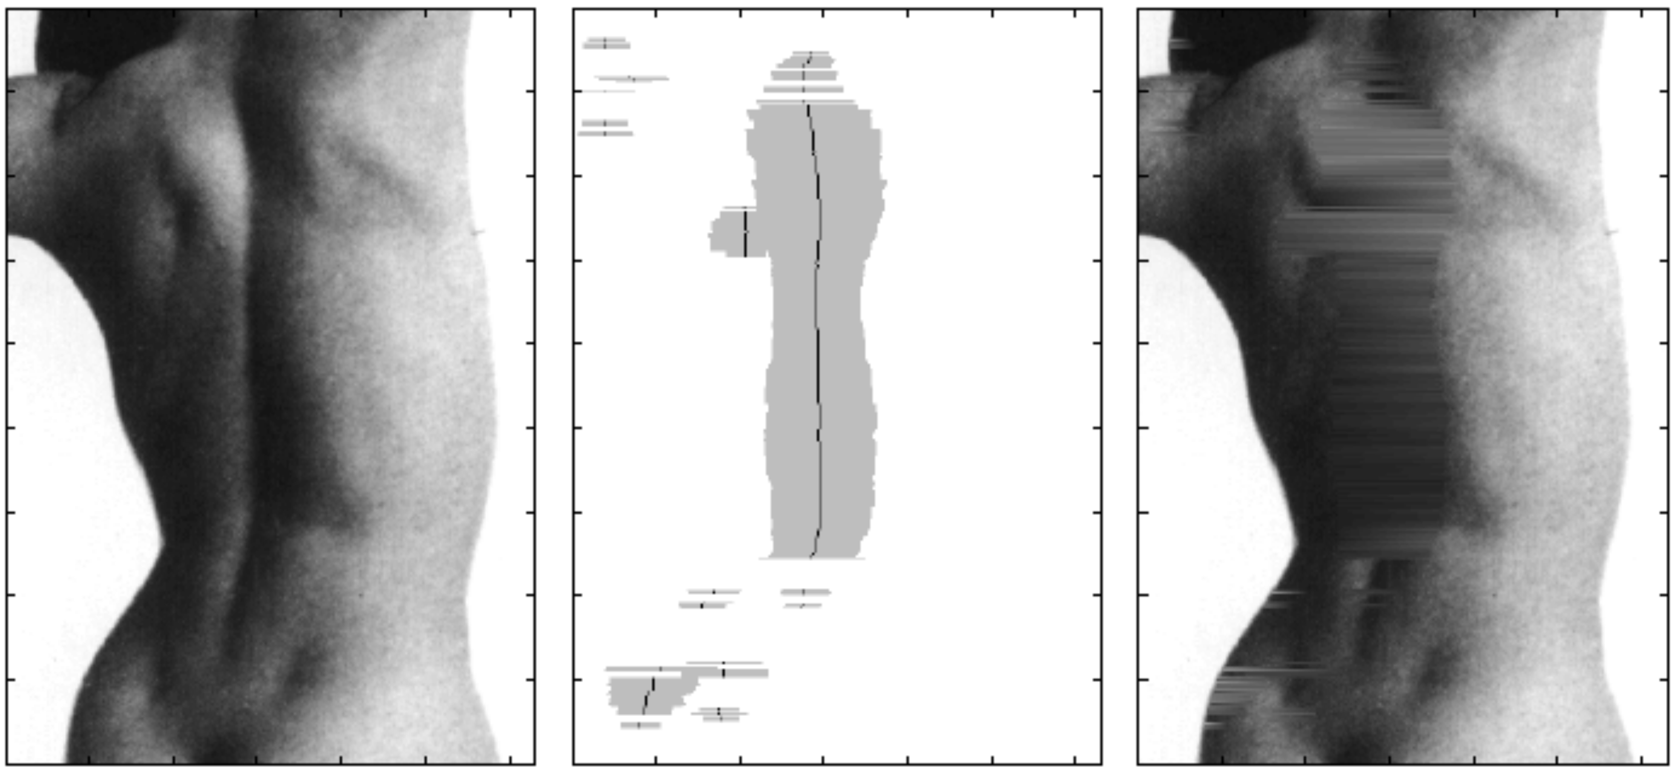

HaddonForsyth 98b

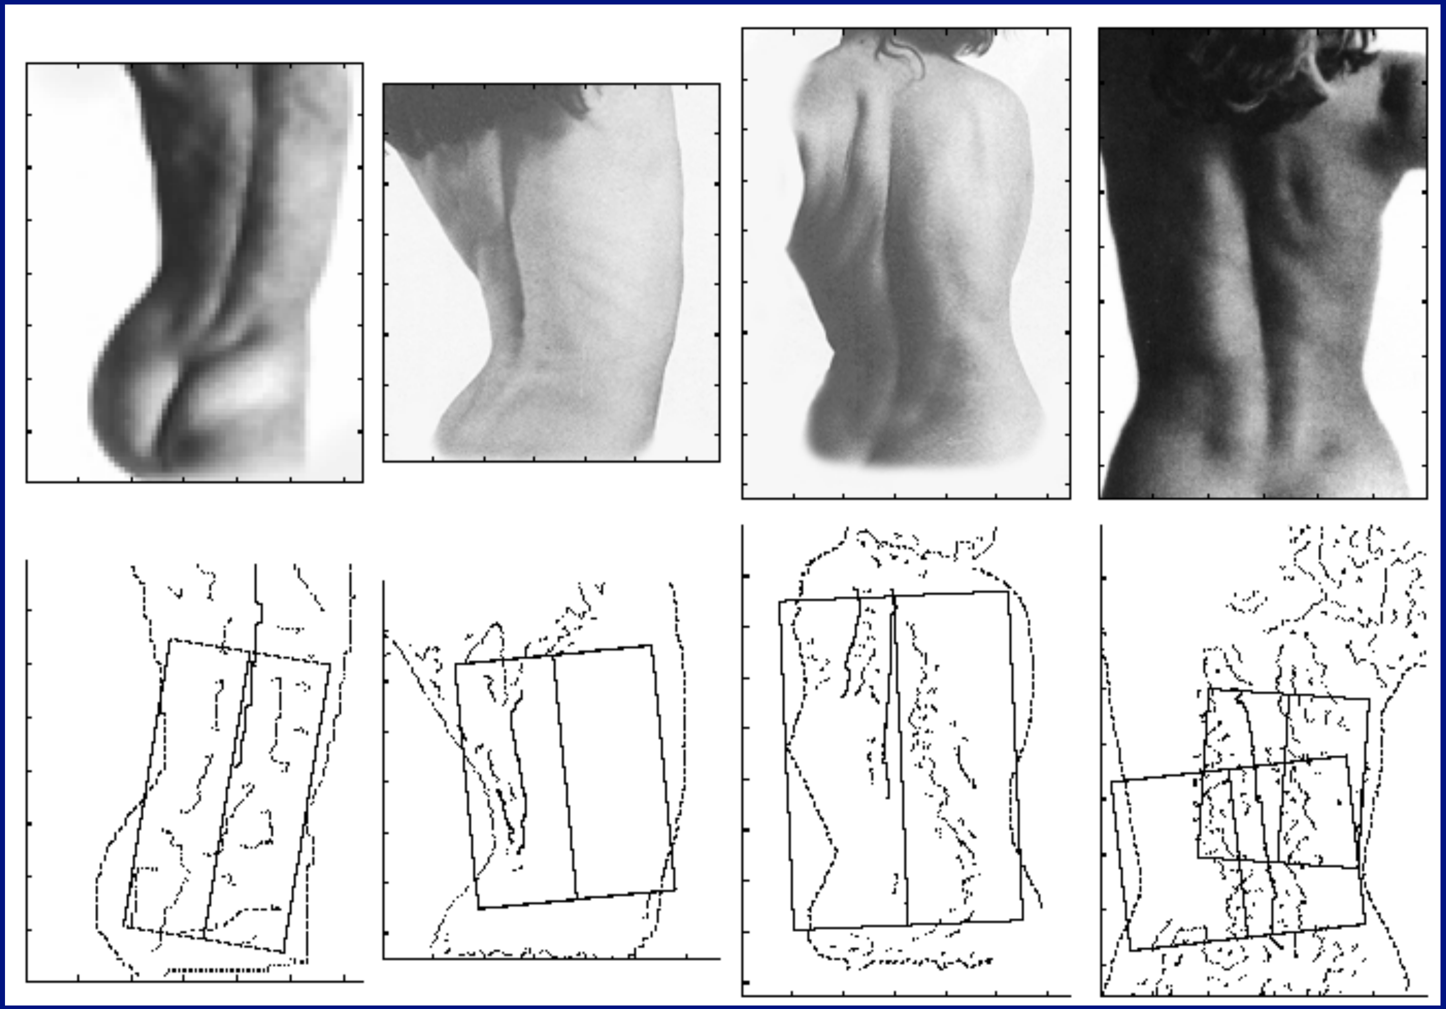

HaddonForsyth 98b






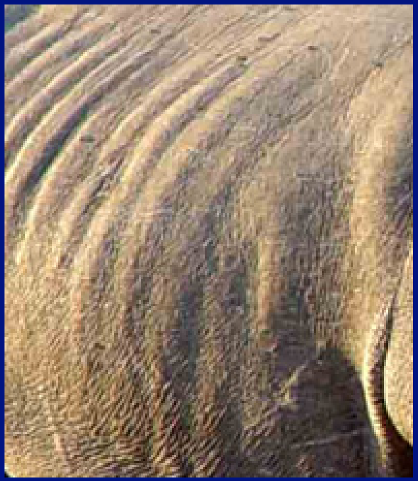












www.naturfoto.cz
© Jan Ševčík


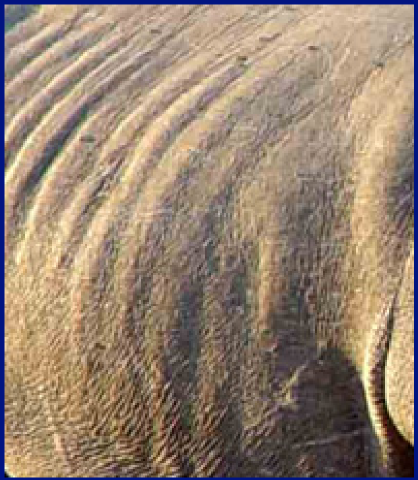



www.naturfoto.cz
© Jan Ševčík


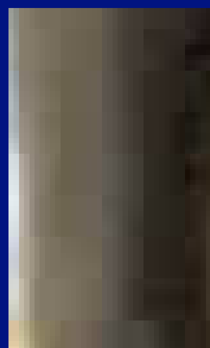


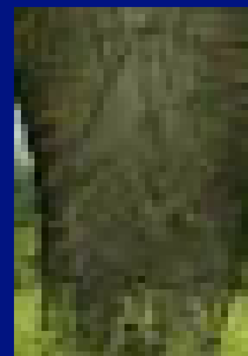

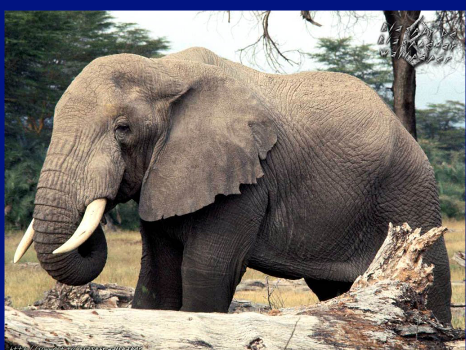

Structure

- Argument:
  - why shading
  - why shading analysis died
  - reasons for hope
- History
  - Classical SFS+Critiques
  - Primitives
- Reconstructions are possible
  - Variable source shading analysis

The Irradiance integral

- Obtain radiosity by
  - summing incoming radiance over all directions

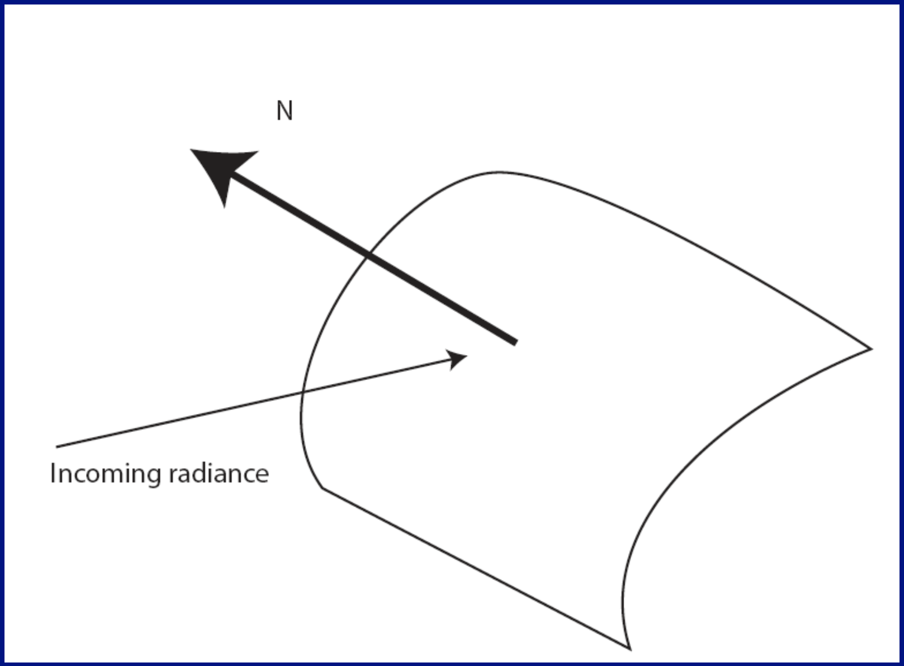

$$B(x,y) = \int_\Omega \rho(x,y;\omega_i) L(x,y;\omega_i) \cos\theta_i d\omega_i$$

The Irradiance Integral

- Classical SFS
  - radiance comes only directly from the luminaire
- Rendering
  - radiance comes from luminaire, reflections from other surfaces

Illumination models and formal solns

Internally generated light

$$B(x) = E(x) + \mathcal{K}B(x)$$

Radiosity

Redistributed light

Internally generated light

One bounce off surfaces

$$B(x) = E(x) + \mathcal{K}E(x) + \mathcal{K}^2E(x) + \mathcal{K}^3E(x) \ldots$$

Gathered from sources

Two bounces off surfaces

$$B(x) = E(x) + \mathcal{K}E(x) + \mathcal{K}(B(x) - E(x))$$

Rendering, Gathering and all that

$$B(x) = E(x) + \mathcal{K}E(x) + \mathcal{K}(B(x) - E(x))$$

$$B(x) = E(x) + \mathcal{K}E(x) + \mathcal{K}(\hat{B}(x) - E(x))$$

- K smoothes
  - phenomena:
    - don’t need a good estimate of B
    - complex angular patterns of radiance are not resolved
      - (Ramamoorthi Hanrahan, 01)
      - useful in photometric stereo (Basri, Jacobs, Kemelmacher 07)

Gathering - II

Quite large, often fast changing (shadows)

$$B(x) = E(x) + \mathcal{K}E(x) + \mathcal{K}(\hat{B}(x) - E(x))$$

Usually zero

Typically slowly changing, usually small

- radiance consists of direct term + indirect term
  - indirect term changes slowly over space
    - irradiance cache (Ward, 88, 92)
    - radiance cache (ArikanForsyth, 04)

Illumination changes slowly over space

Radiance Cache Samples

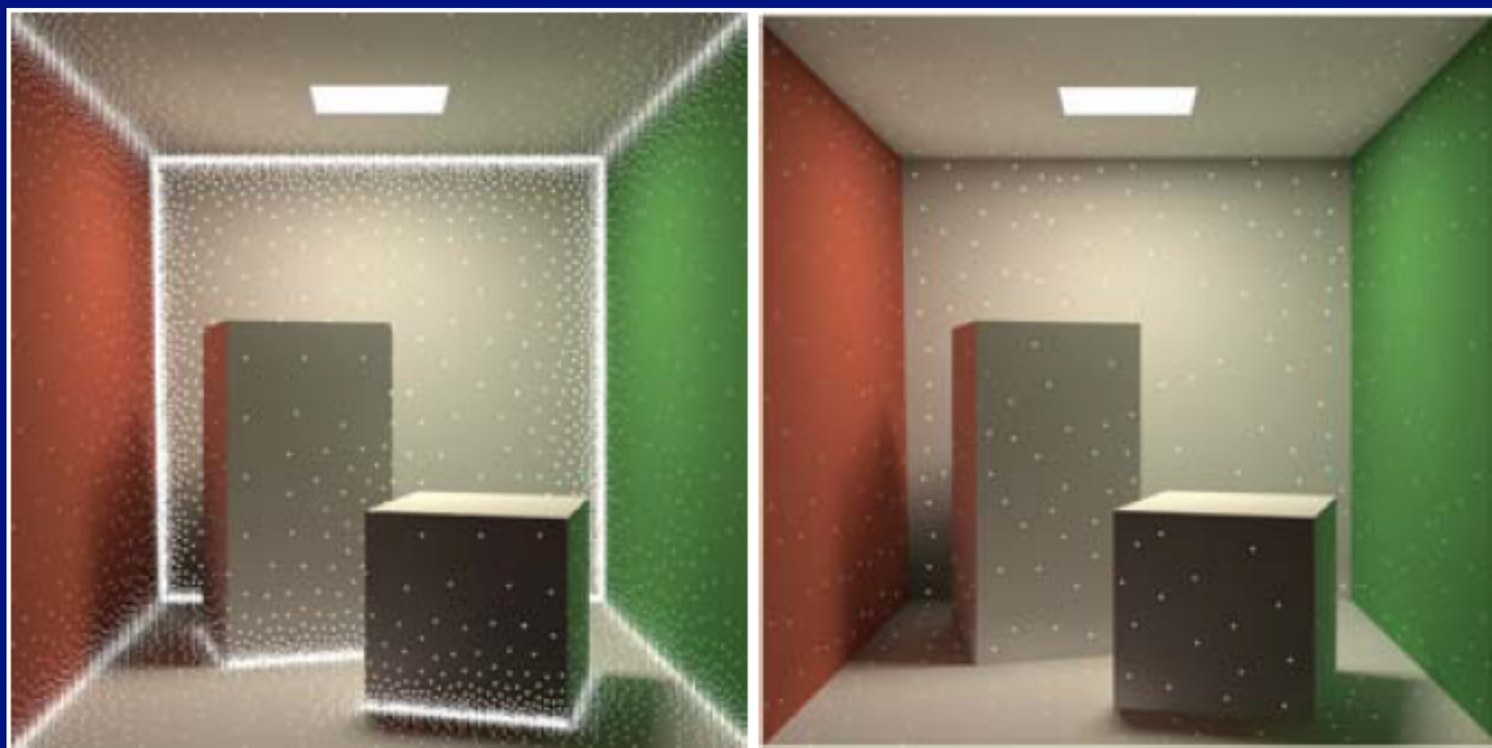

Irradiance Cache Samples

Figure from Arikan Forsyth 05

The effective source

$$R(p, q; \mathbf{S}_e(x, y)) = I(x, y)$$

- A spatially varying source
  - that produces the right answer from the reflectance map
- Properties
  - not very different from ideal source
  - difference changes slowly over space

Variable Source Shading Analysis

Minimize

Slow change in effective source

Effective sources similar to source

$$\theta_1 \sum_{i \in Sources} \int_\Omega \parallel \nabla \mathbf{S}_e^{(i)}(x,y) \parallel^2 dA \quad + \quad \theta_2 \sum_{i \in Sources} \int_\Omega \parallel \mathbf{S}_e^{(i)}(x,y) - \mathbf{S} \parallel^2 dA +$$

$$\theta_3 \int_\Omega (f_{xx} + f_{yy})^2 dA \quad + \quad \theta_4 \left(\int_\Omega dA_s - A_0 \right)^2$$

No free creases

Extra area is expensive

Subject to:

$$R(p,q;\mathbf{S}_e^{(i)}(x,y)) = I(x,y)$$

Boundary conditions

Variable source shading analysis

- Solution always exists
  - if boundary conditions are consistent
- Arbitrary (consistent) boundary conditions OK
- Can do 0, 1, 2.... sources
- Area regularizer is very helpful
- Somewhat stabler problem if we substitute:

$$\mathbf{S}_e^{(i)}(x, y) = g_i(x, y)\mathbf{S}^{(i)}$$

Local shading model


Physically realistic shading

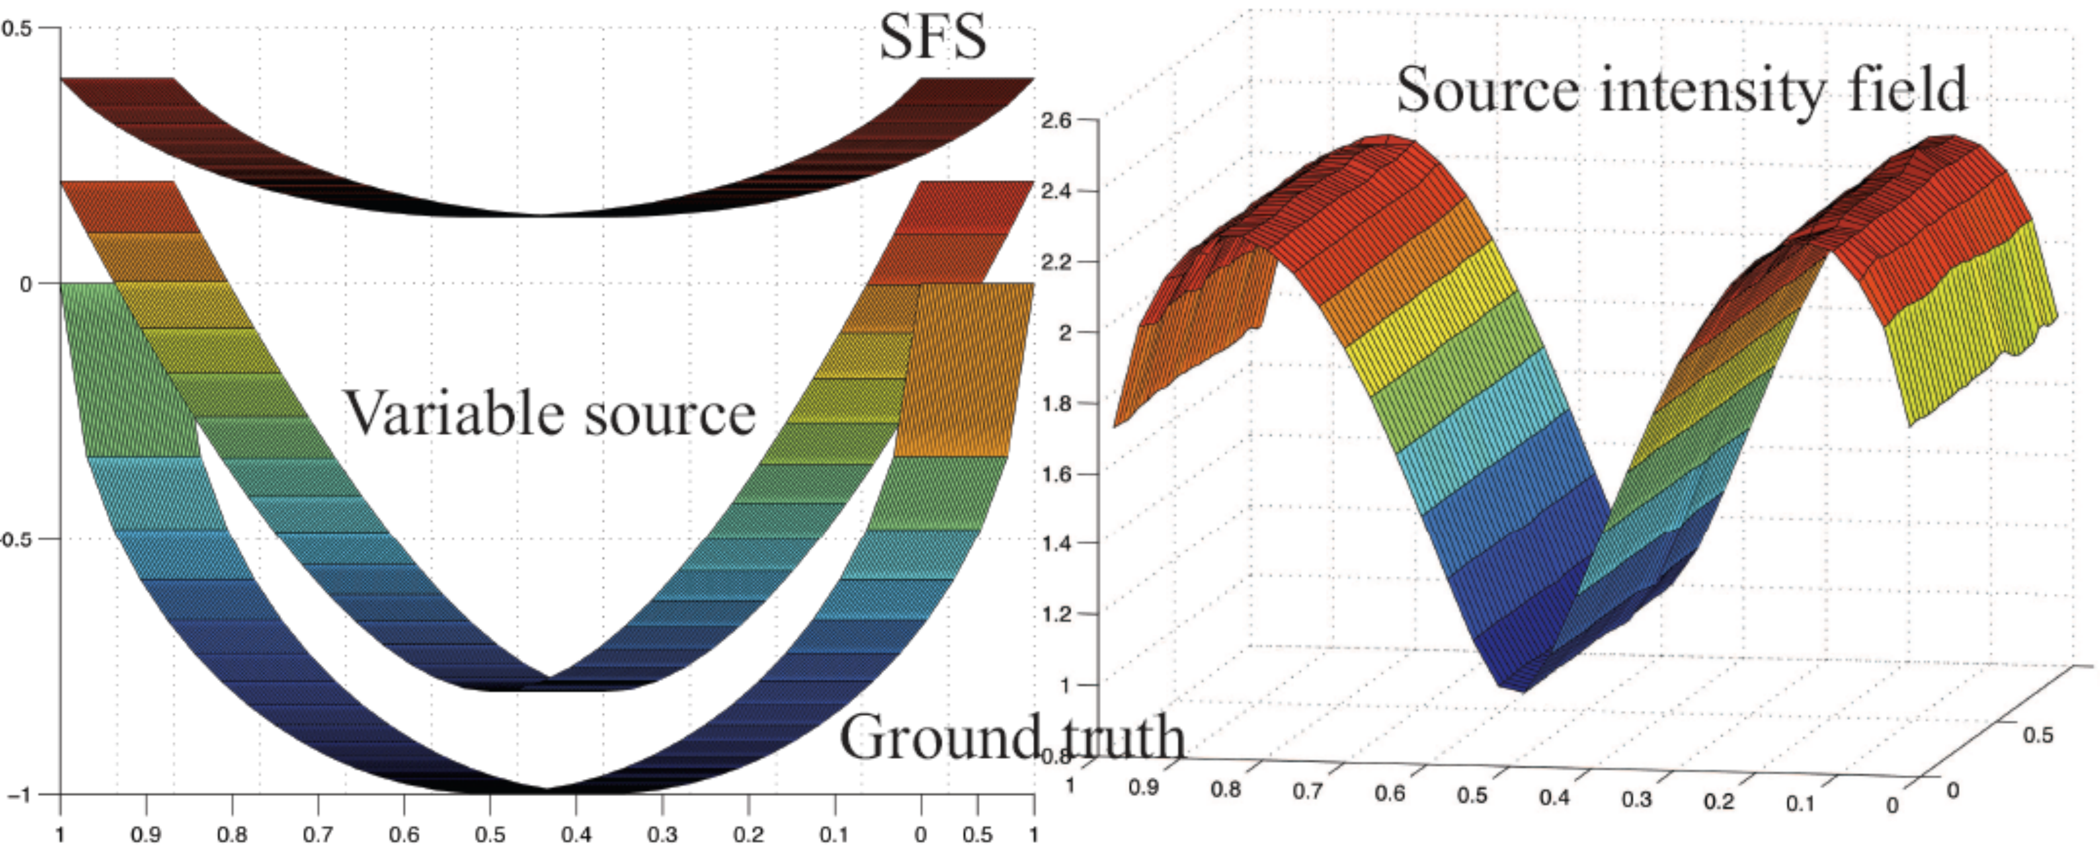

SFS
Variable source
Ground truth
Source intensity field

Figures 1a, b of Koenderink, “Pictorial Relief”, 98

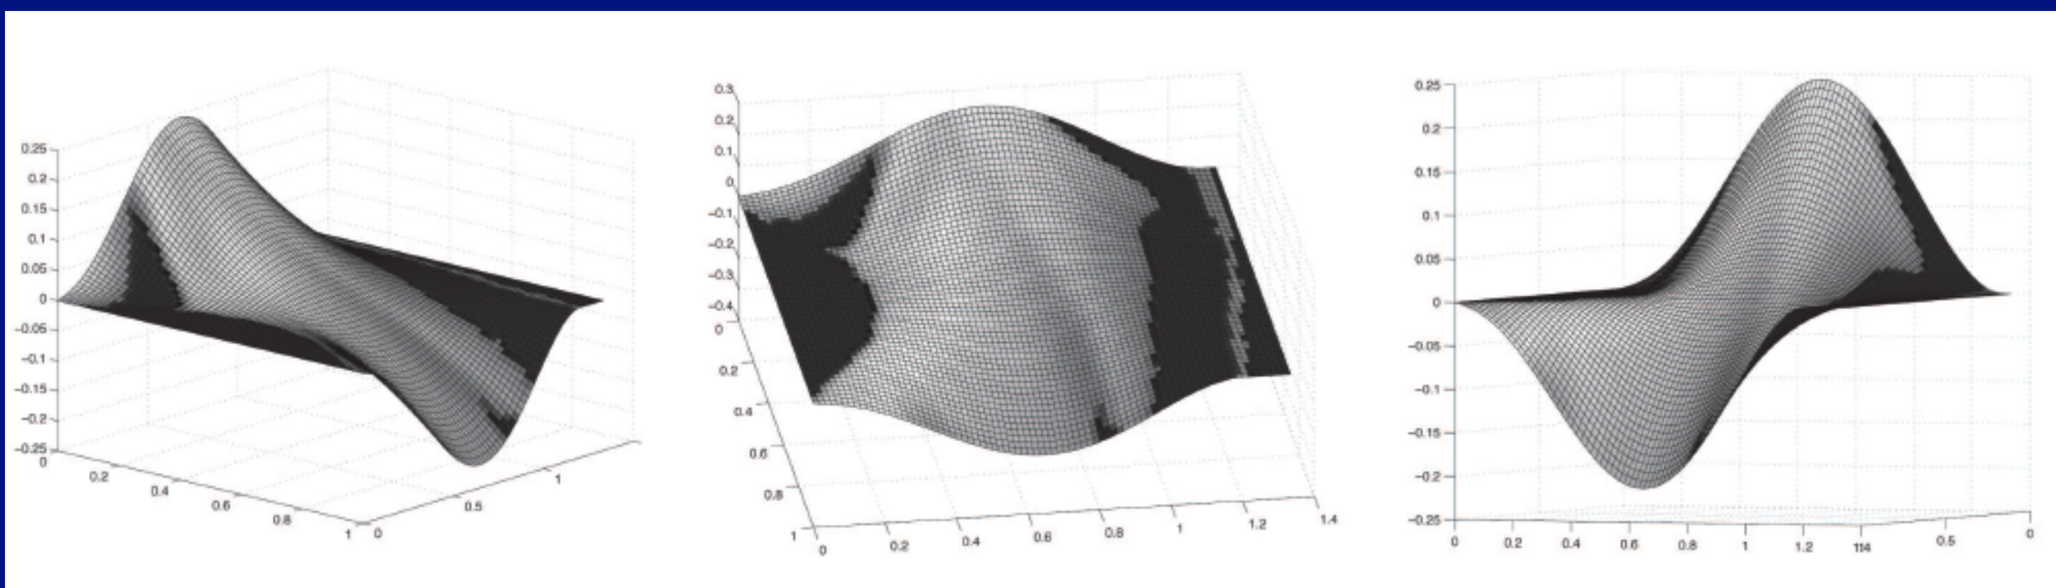

No shading (this isn’t unique, but gives some idea of what bc’s do)

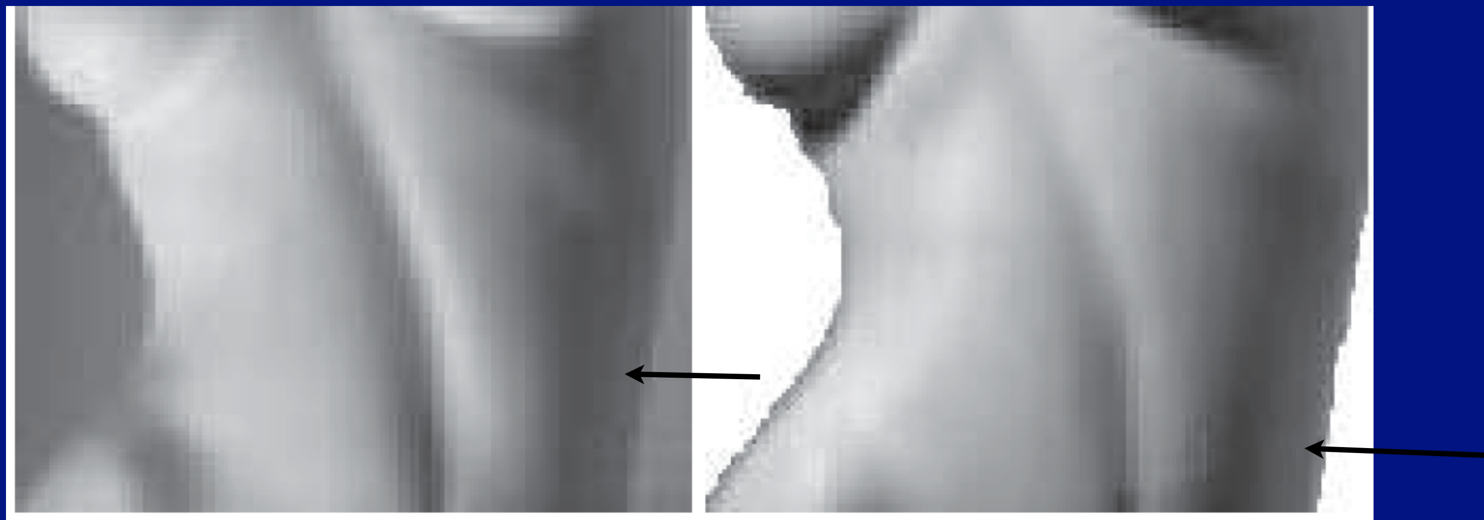

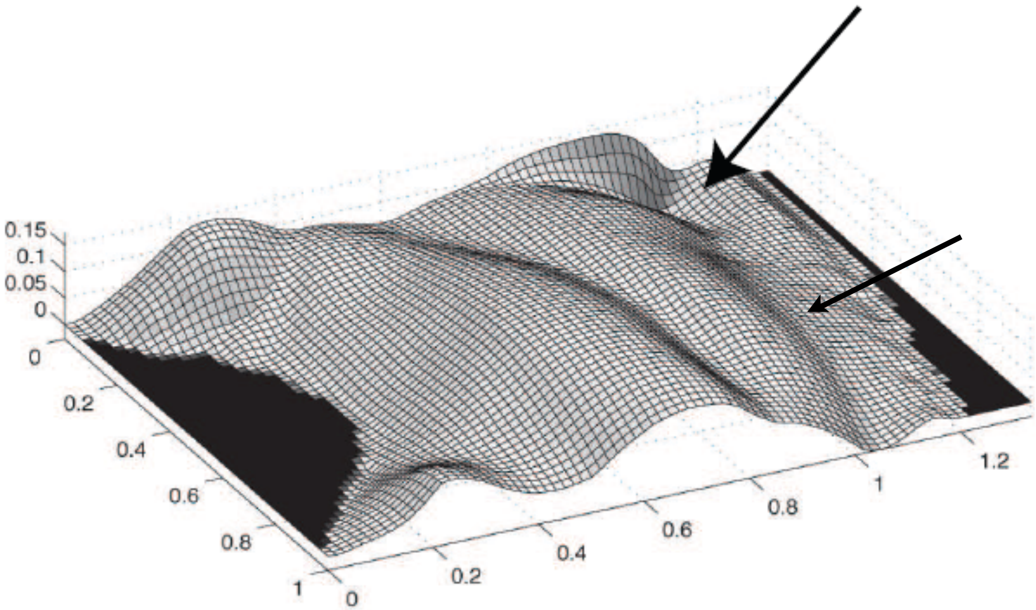

2 Source reconstruction

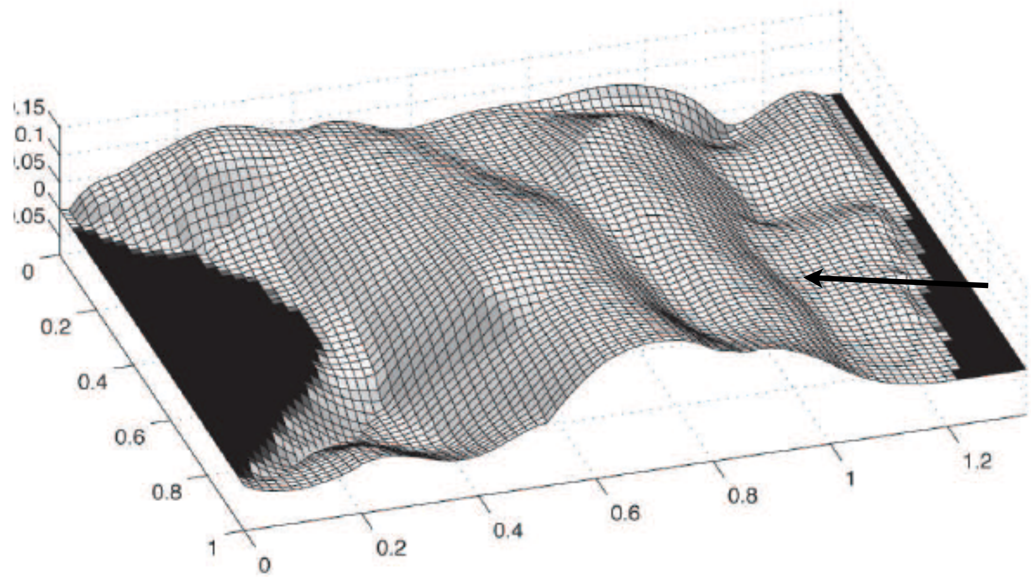

1 Source reconstruction

1-Source vs 2 Sources

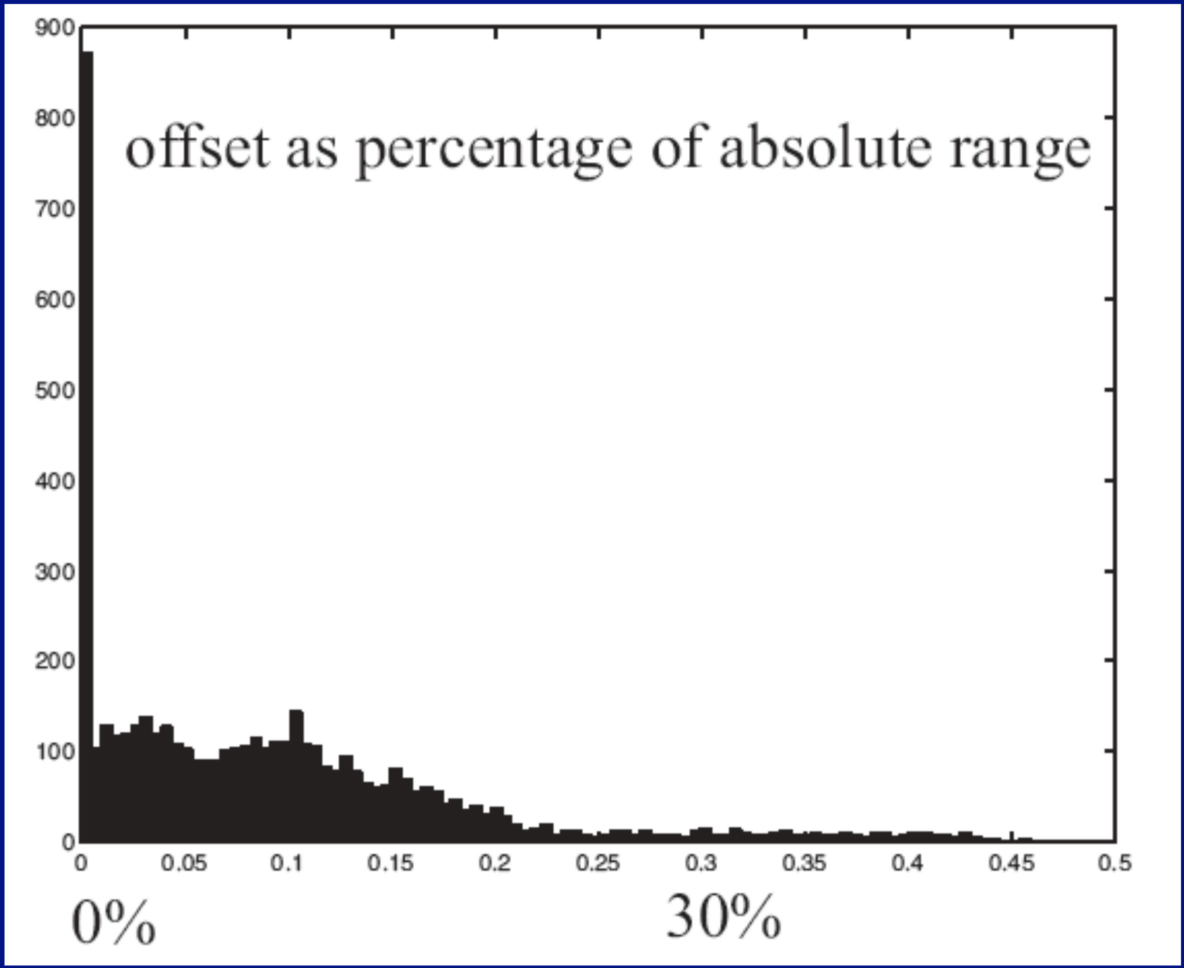

Masked image

Albedo
(inferred from photometric stereo and provided)

Shading image

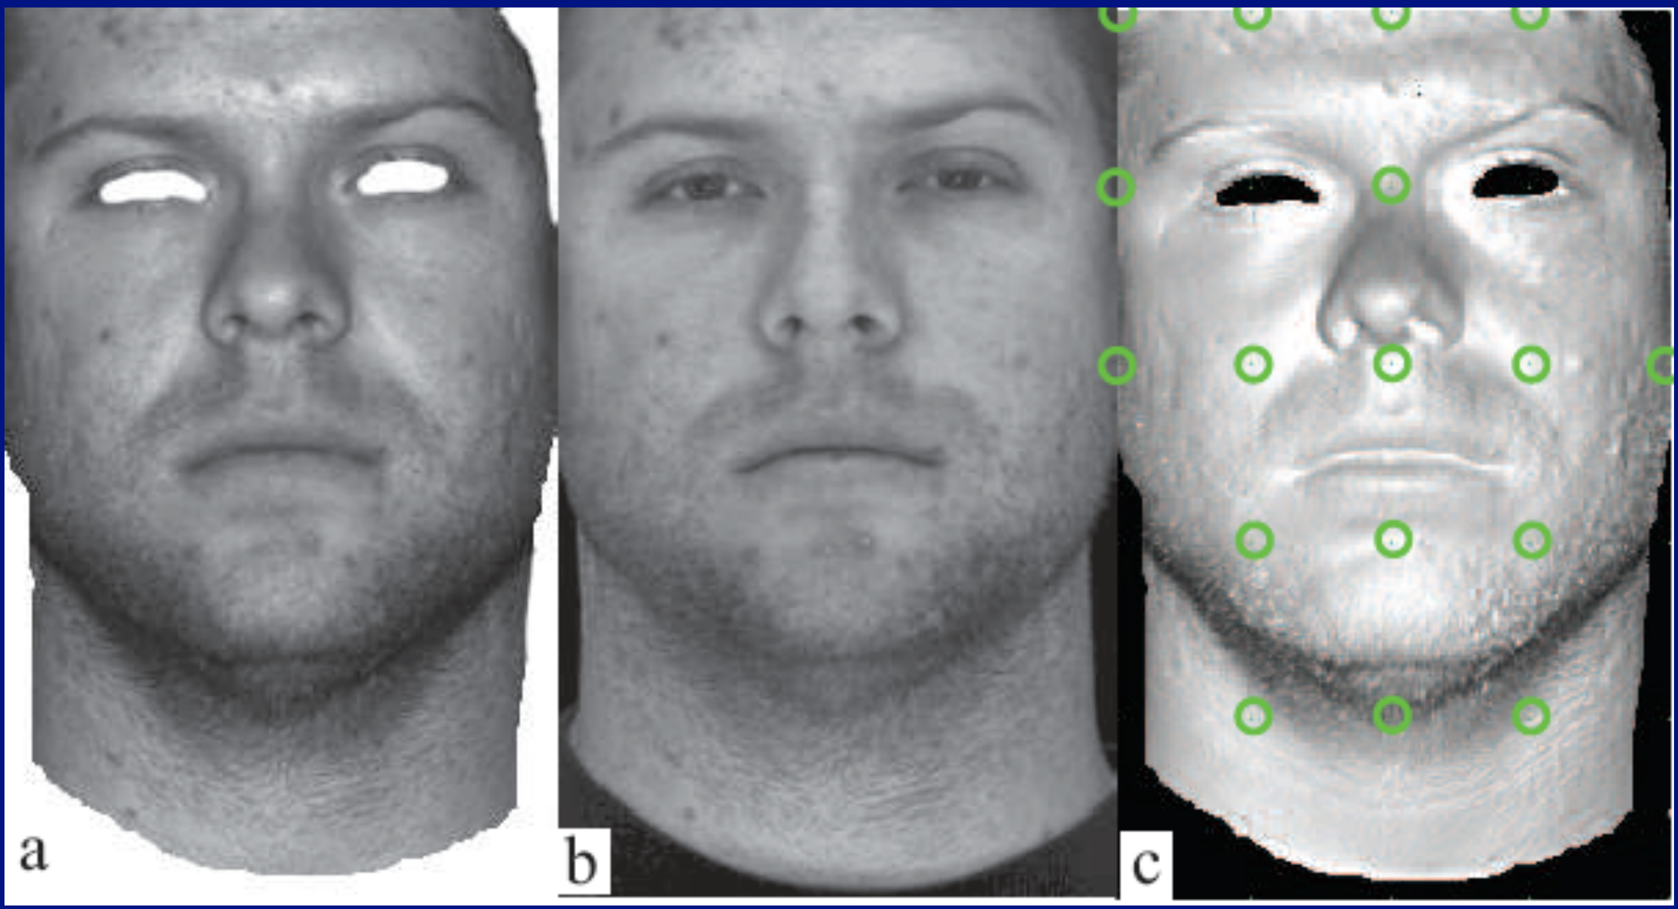

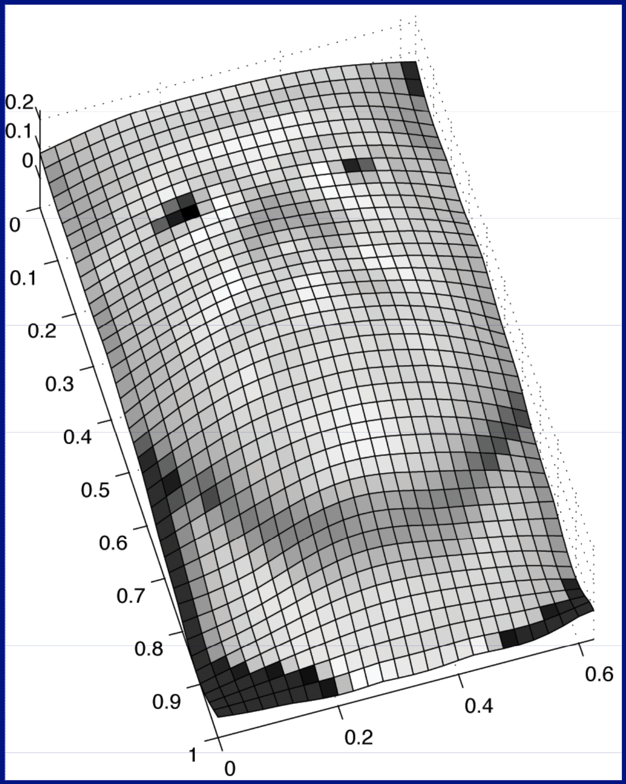

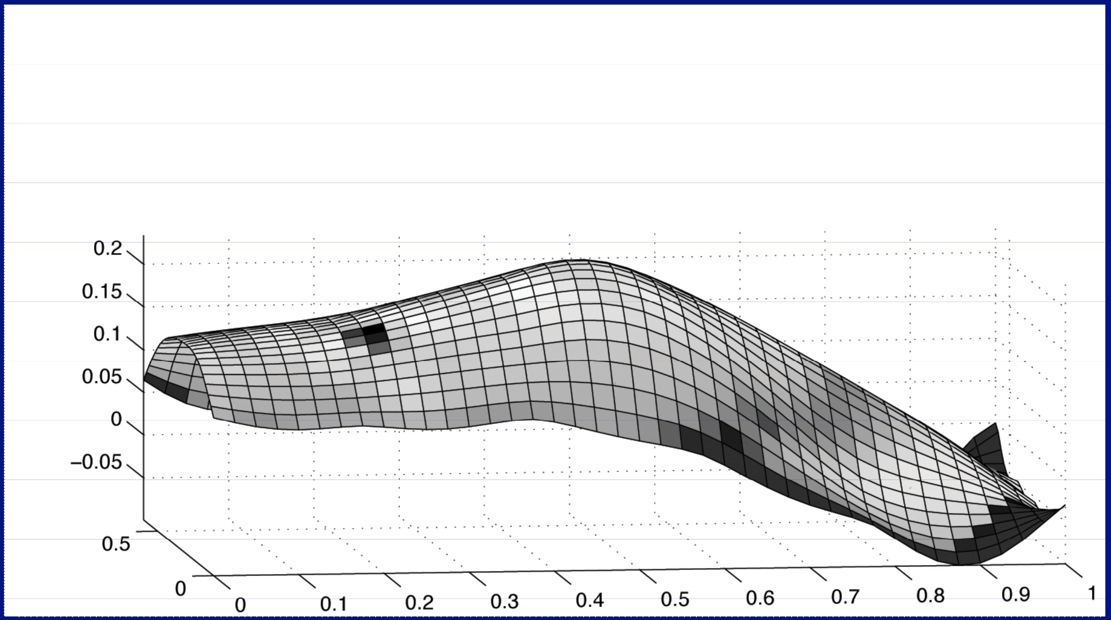

Without shading

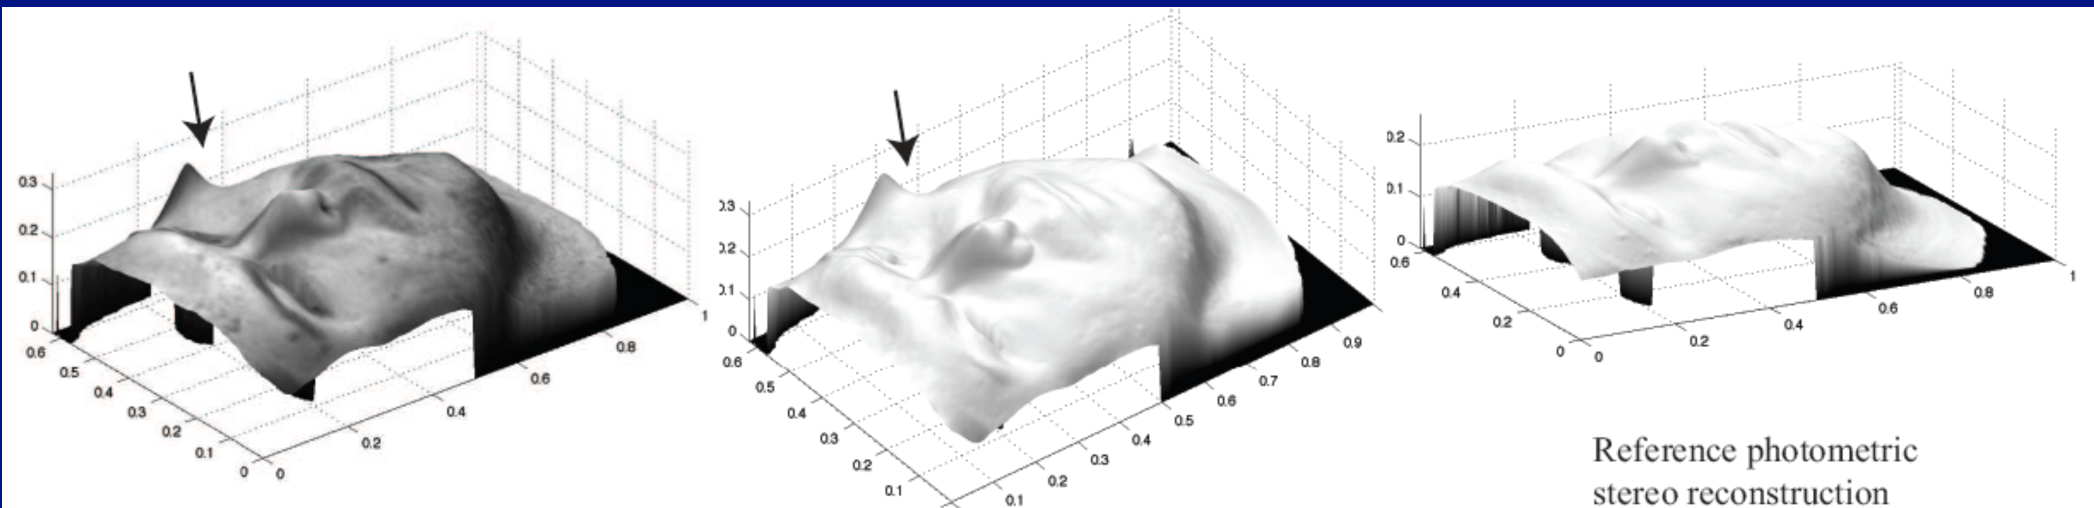

Single source reconstruction with albedo

Single source reconstruction without albedo

Reference photometric stereo reconstruction

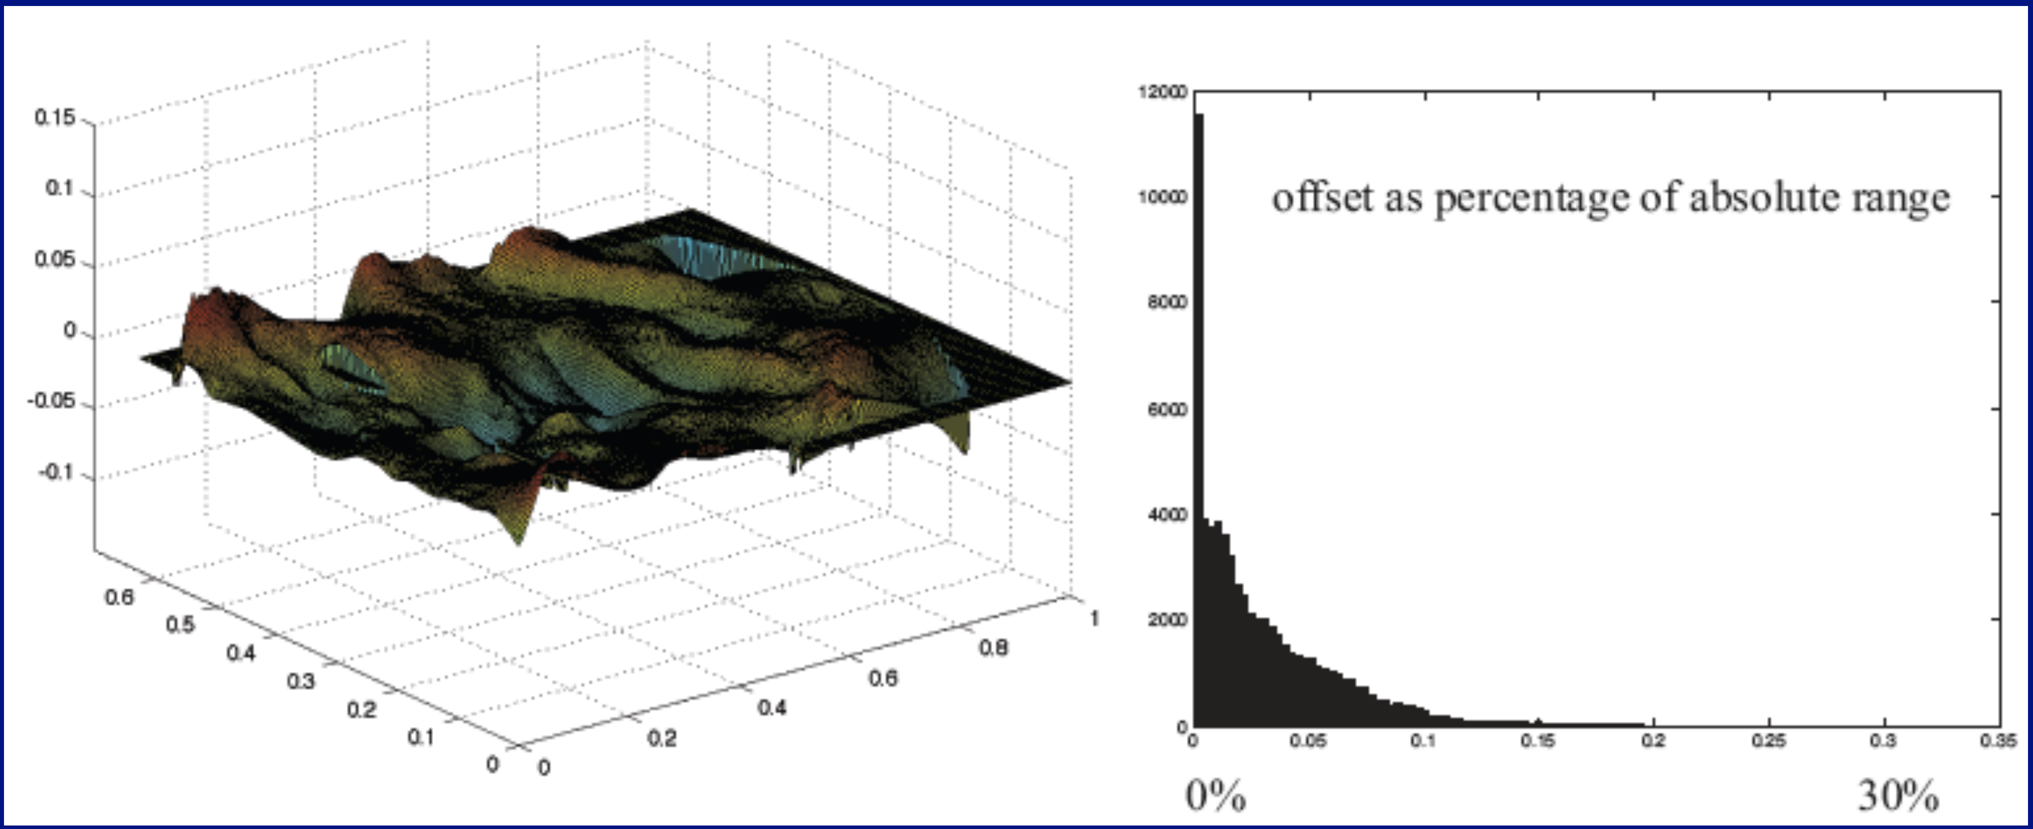

Single source face against reference photometric stereo reconstruction

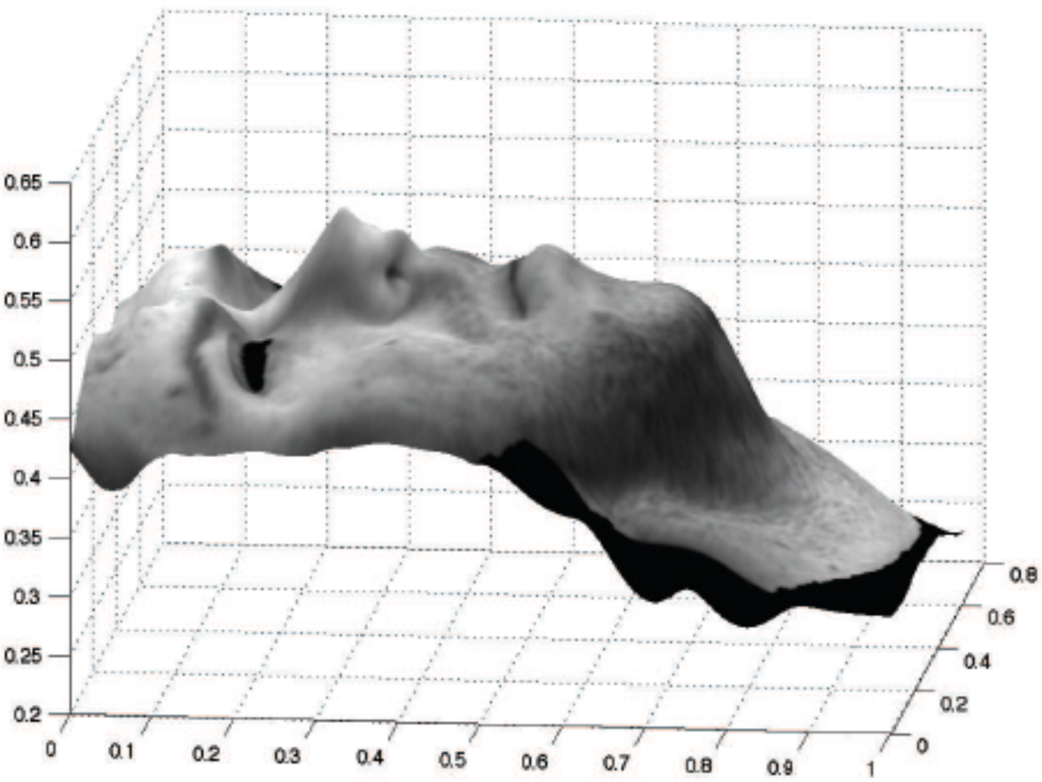

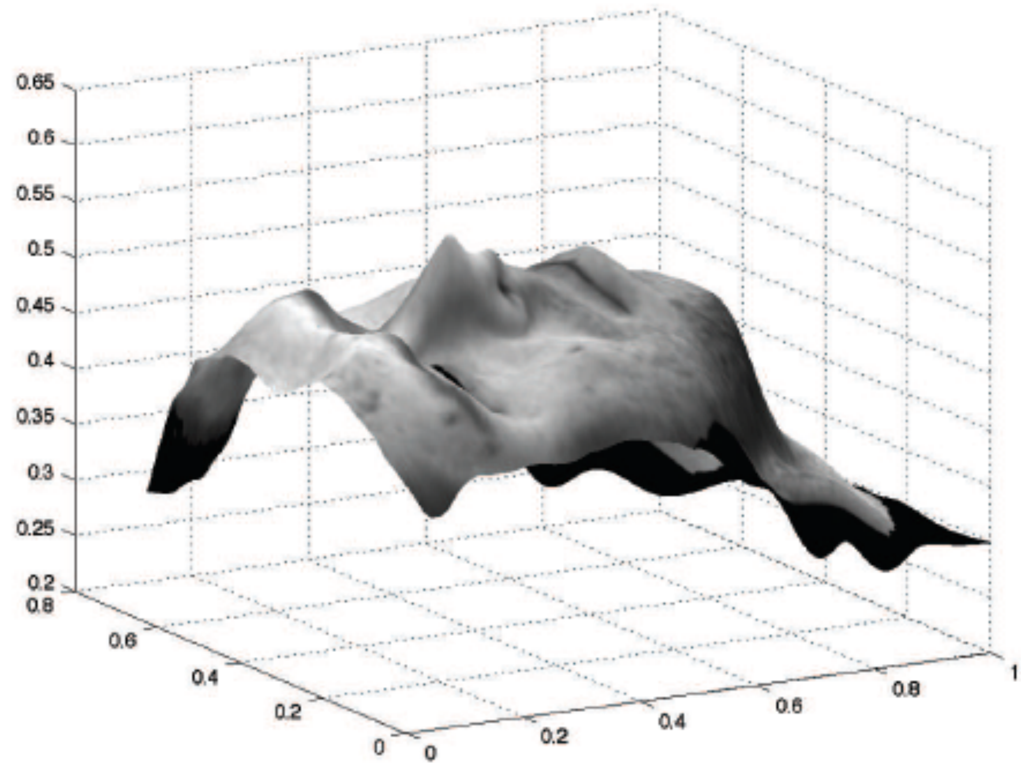

320x200 representation:
single source
256000 variables
640 depth constraints (32x20 grid) some masked
Note bump on nose - specularity

Snow field, from flickr, by tecnhicolouryawp

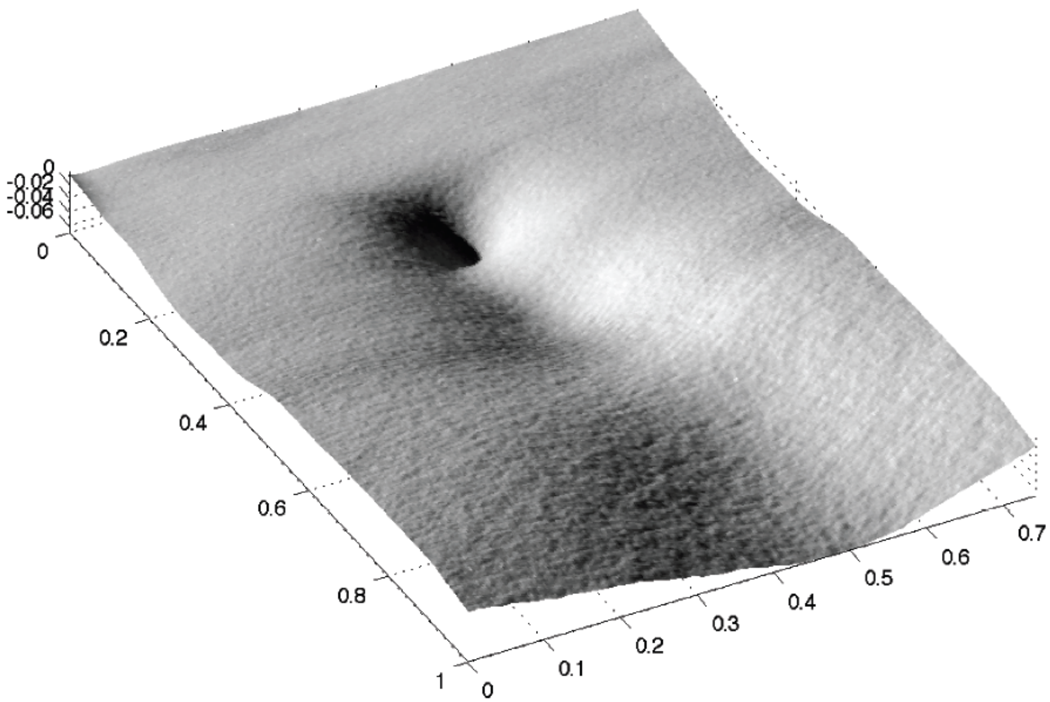

0
-0.02
-0.04
-0.06
0
0.2
0.4
0.6
0.8
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7

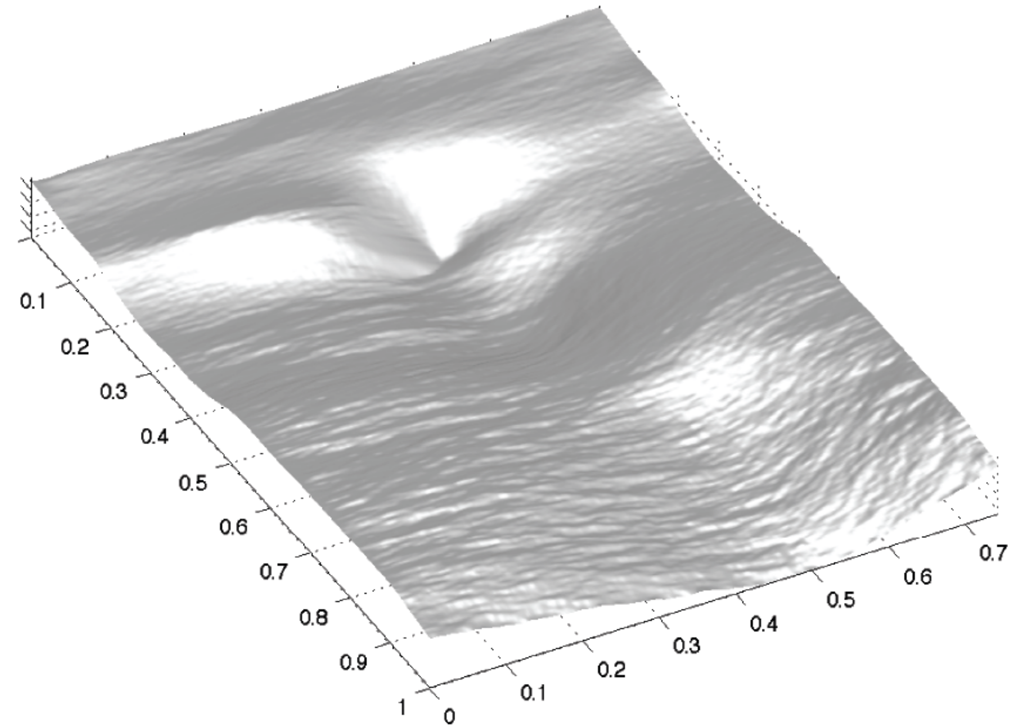

0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7

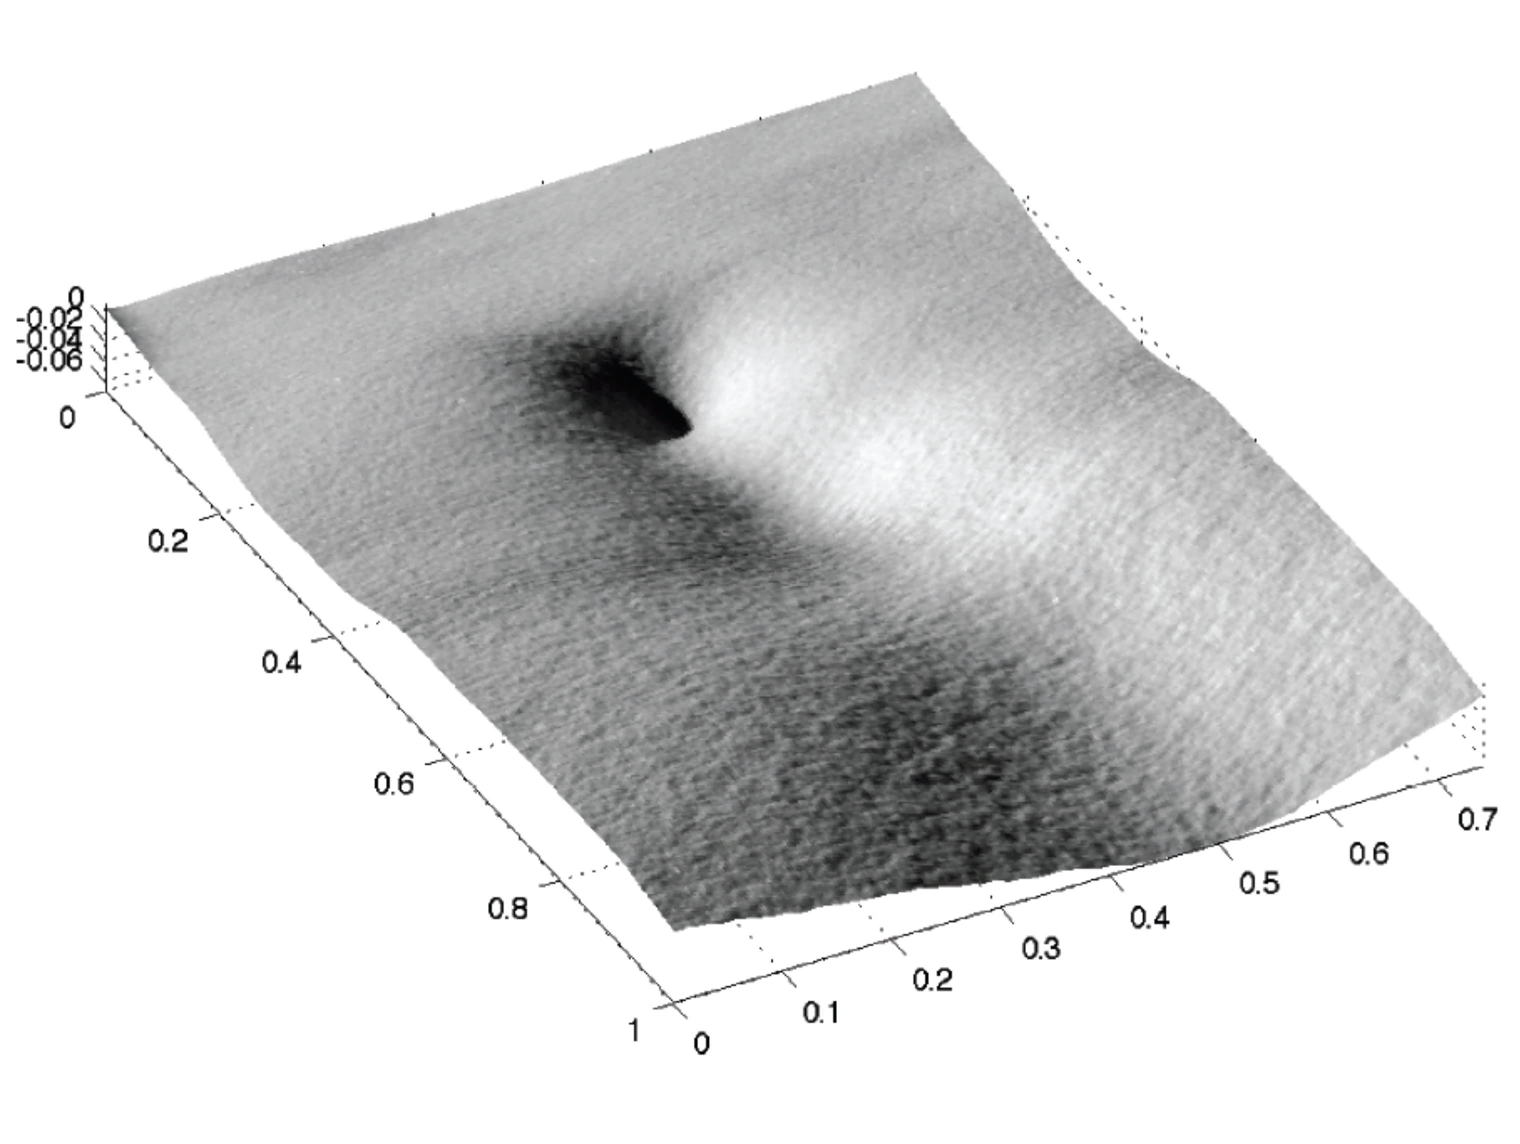

0
-0.02
-0.04
-0.06
0
0.2
0.4
0.6
0.8
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7

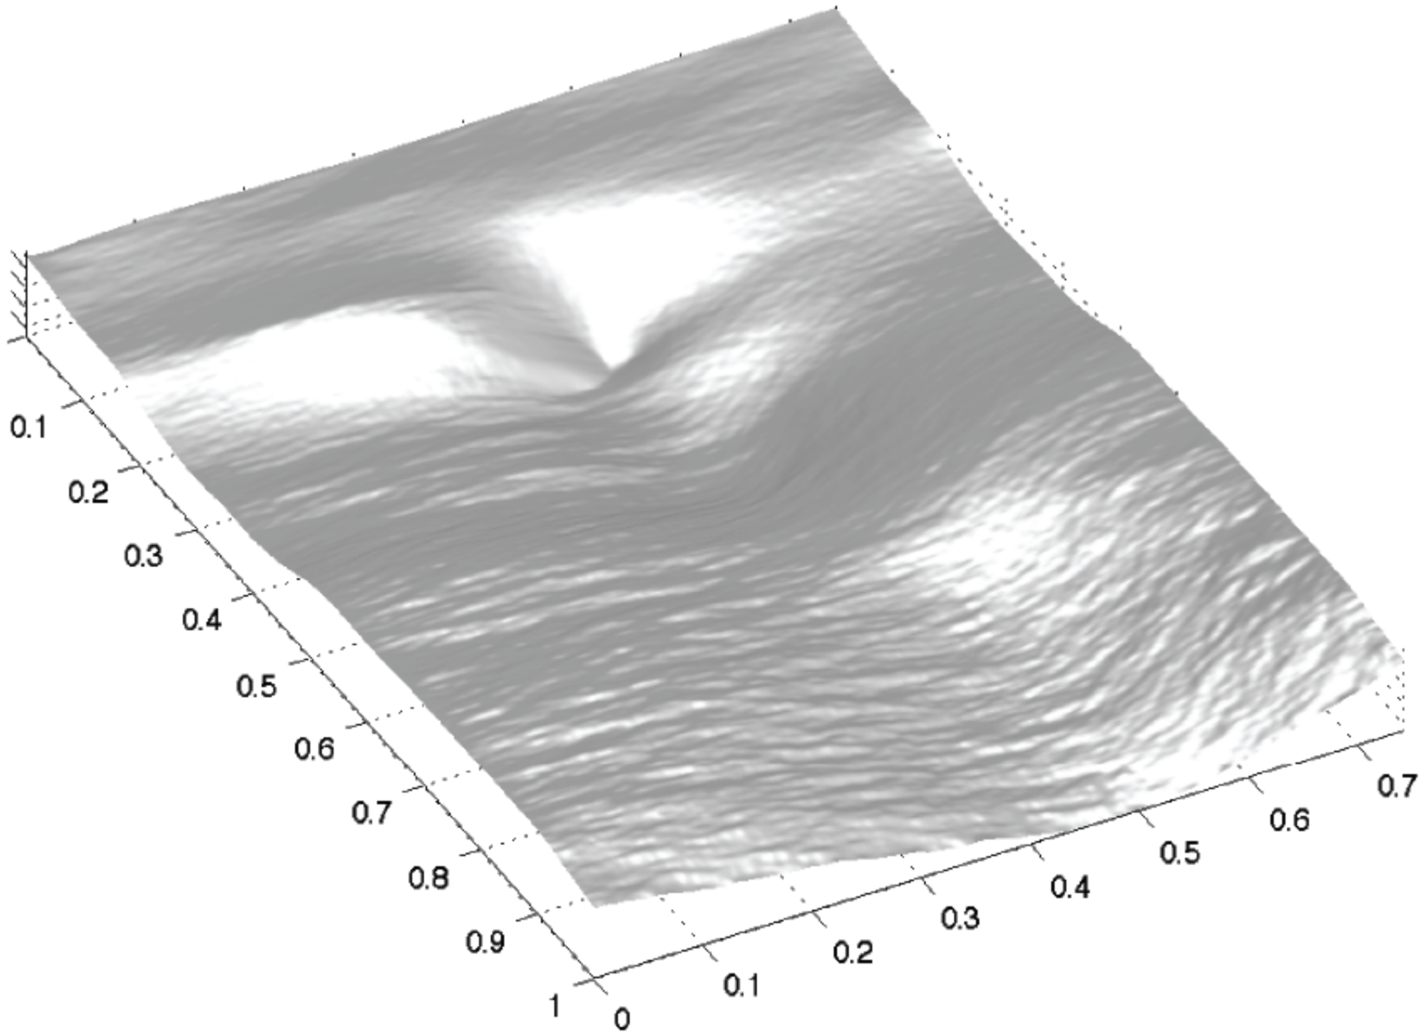

0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7

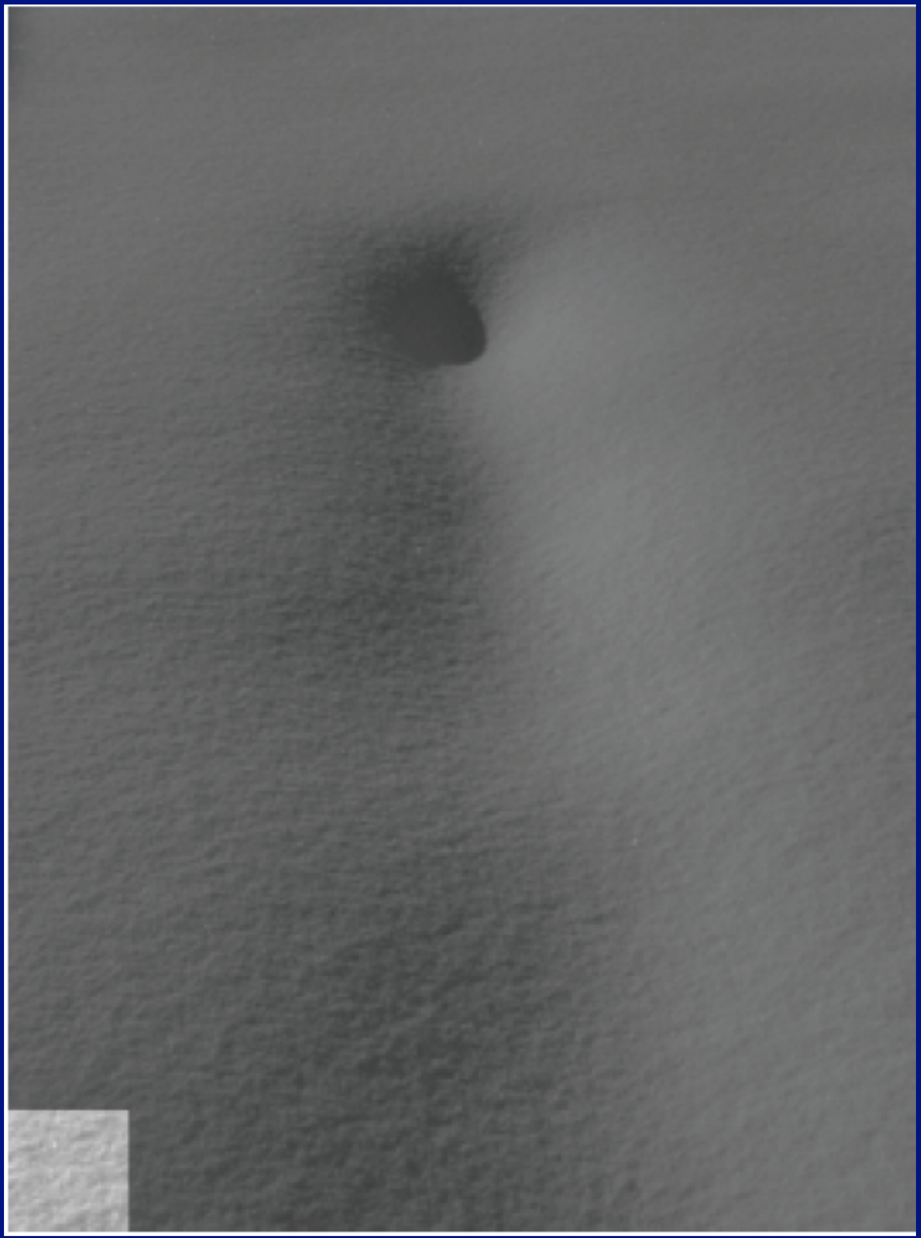

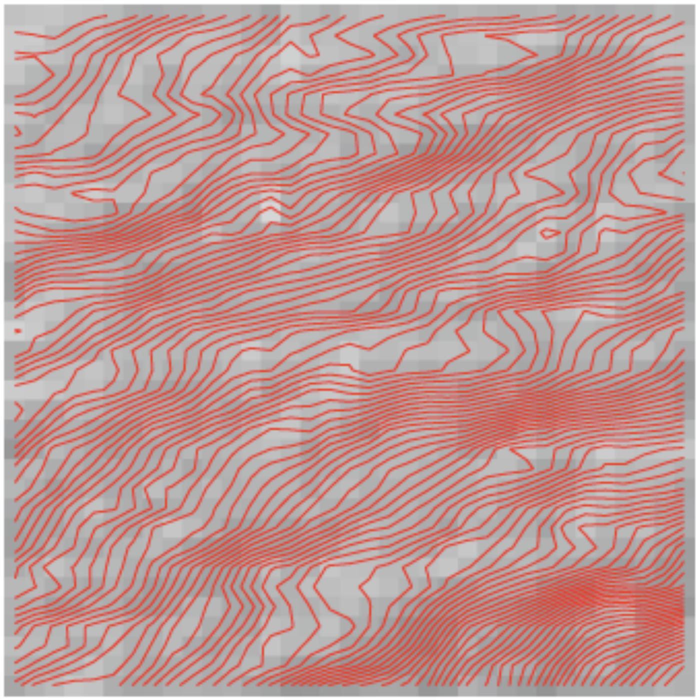

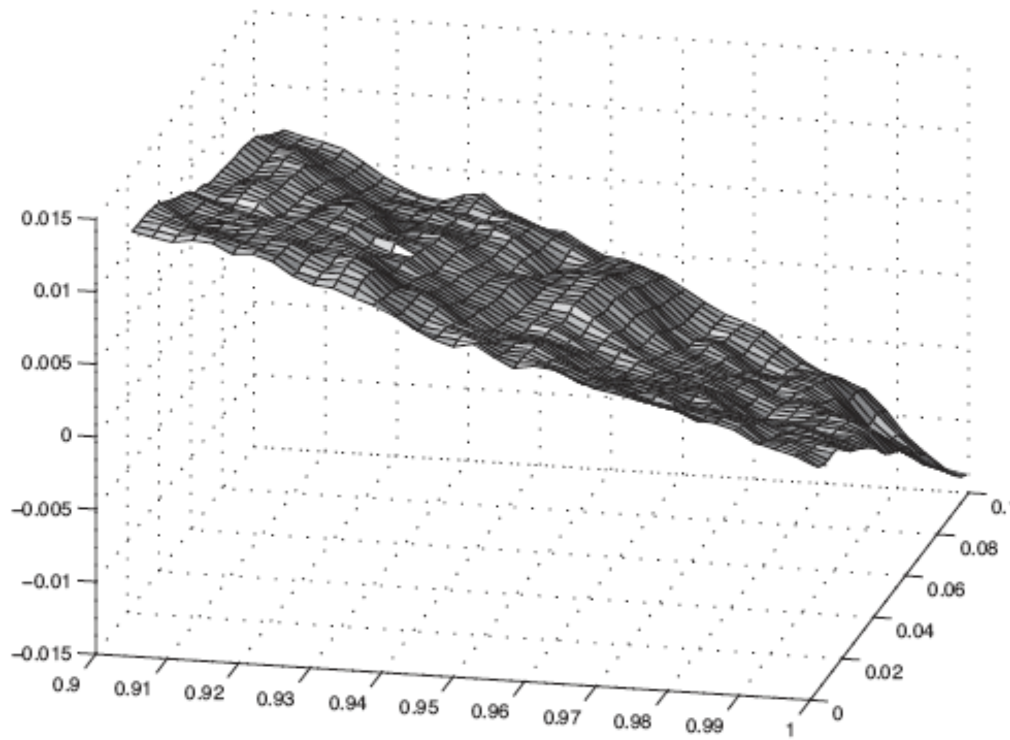
0.015
0.01
0.005
0
−0.005
−0.01
−0.015
0.9
0.91
0.92
0.93
0.94
0.95
0.96
0.97
0.98
0.99
1
0
0.02
0.04
0.06
0.08
0.1

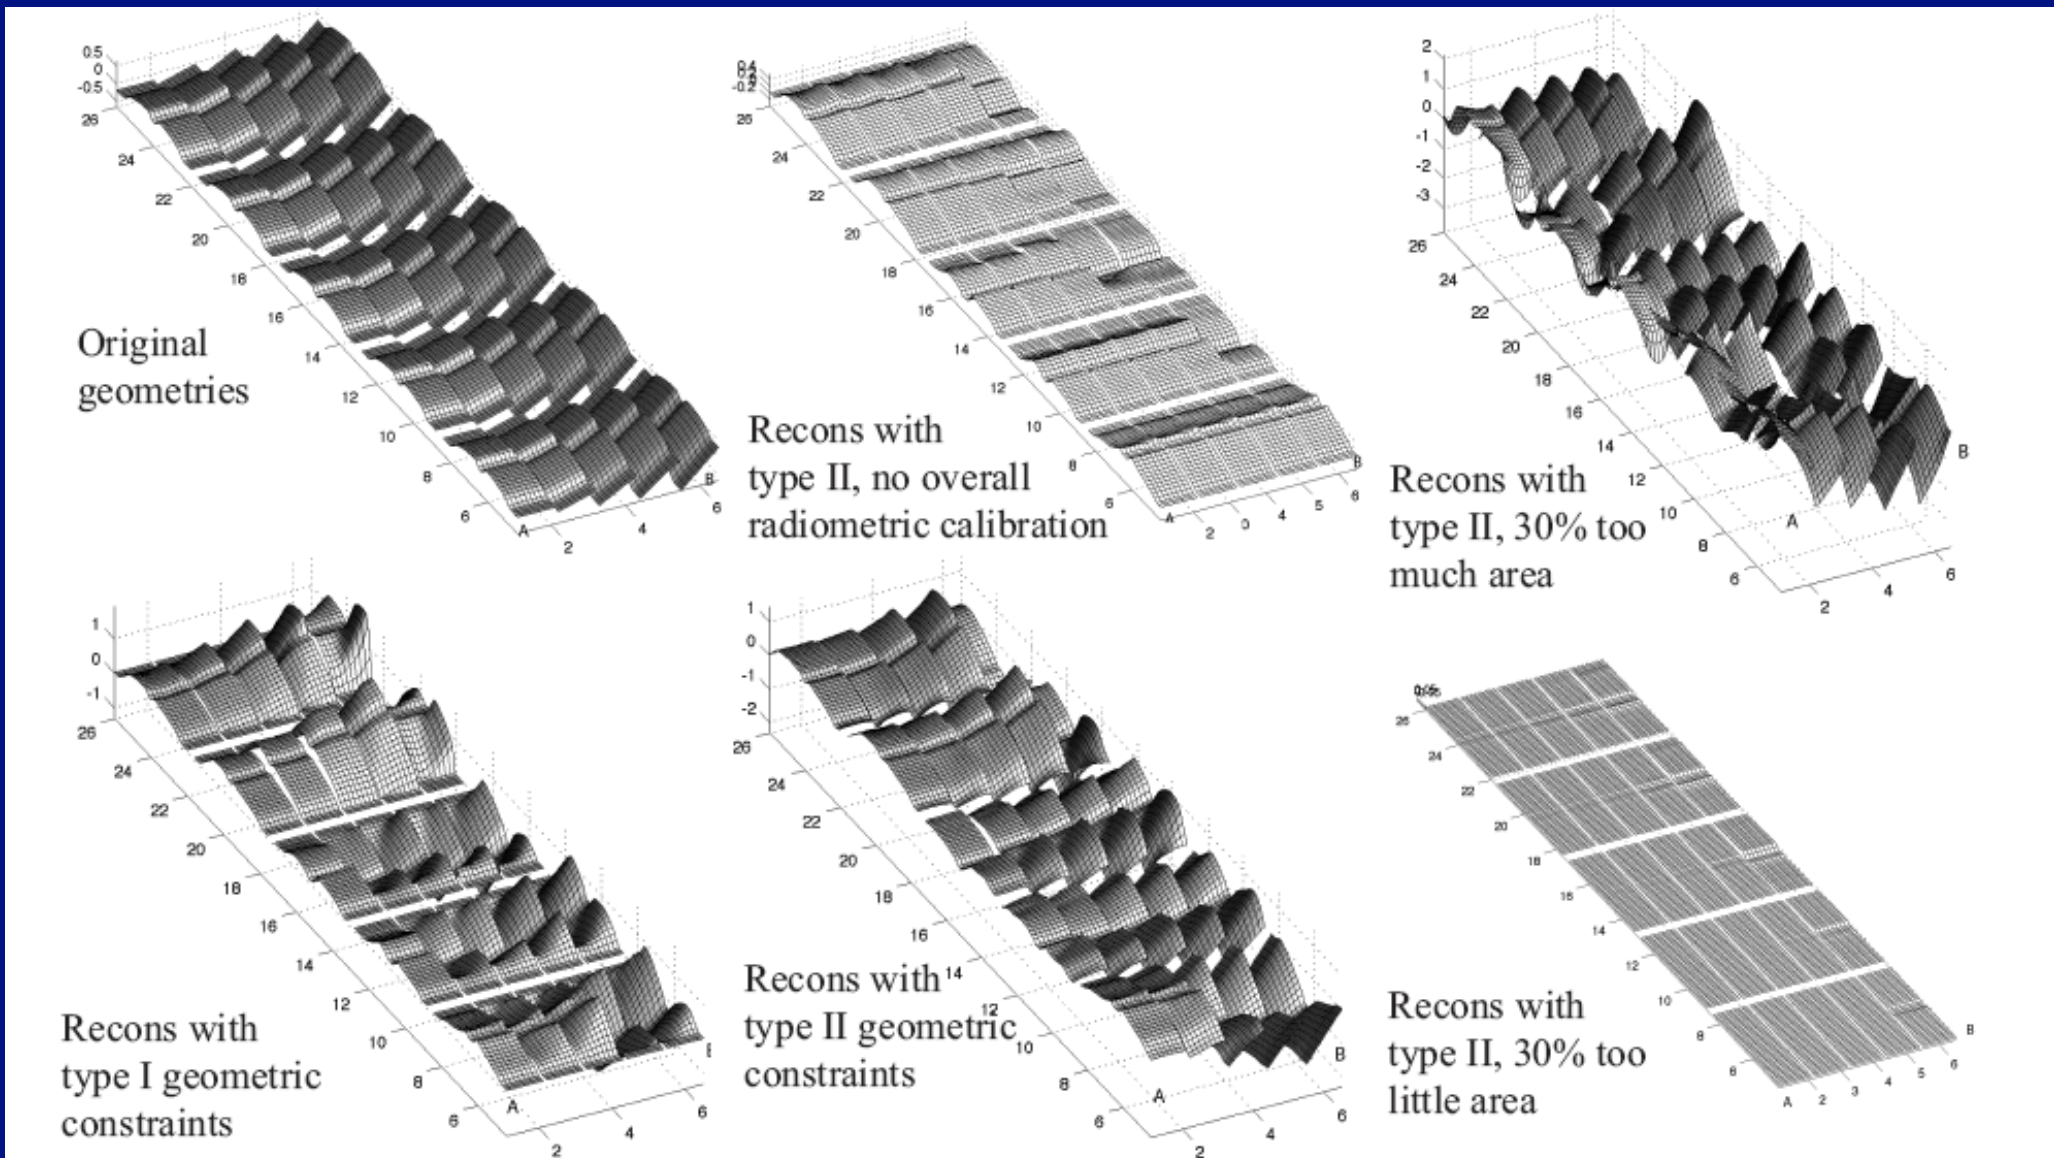
Original geometries
Recons with type II, no overall radiometric calibration
Recons with type II, 30% too much area
Recons with type I geometric constraints
Recons with type II geometric constraints
Recons with type II, 30% too little area

Further matters

Fine scale shadows are a cue to projected illumination direction - this is local (Koenderink).

Can this be fused with brightness reasoning?

Further matters

Humans seem to be able to reason about light in in volume - how bright would something be if we put it here?

Can we mimic this?

Important points

- There are features which exist over spatial domains
  - at object length scales
- Usable notion of primitive essential
  - to handle unknown objects
- The visual world is very rich
  - cue opportunism is essential for both reconstruction and recognition
- Q: Should we reconstruct for recognition
  - A: ?