Multi-view shape reconstruction

Shape reconstruction

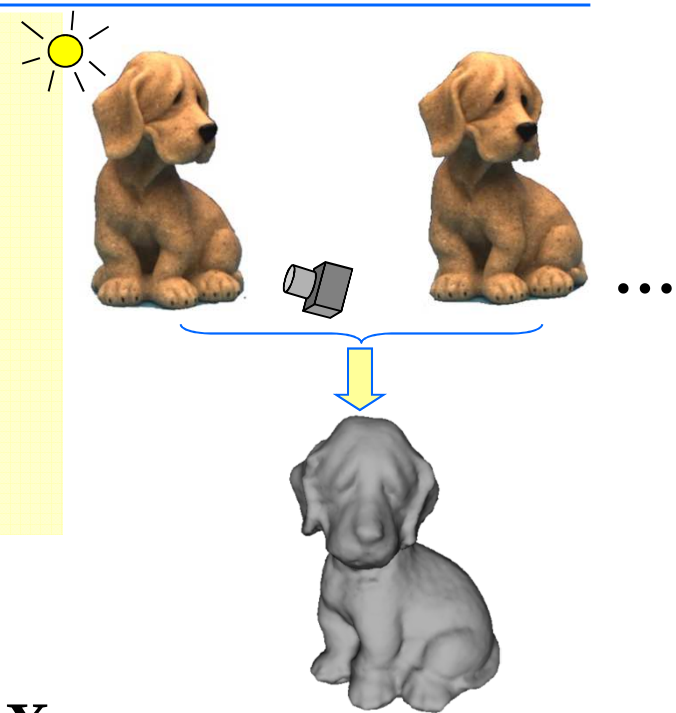

Given

- A set of images (views) of an object / scene
- Camera calibration information
- [Light calibration information]

Find the surface that best agrees with the input images.

Approach:

- chose a surface representation $S, \mathbf{X}$
- define a photo-consistency function $\phi(\mathbf{X})$

  [in practice photo-consistency+regularization]

- solve the following minimization $\min_S \int_{X \in S} \phi(X) dX$

Photo-consistency function(al)

$\phi(X)$

Based on image cues (shading, stereo, silhouettes, …)

Extension SFS/PS to multi-view:

Needs camera/light calibration !

Move camera

SFS

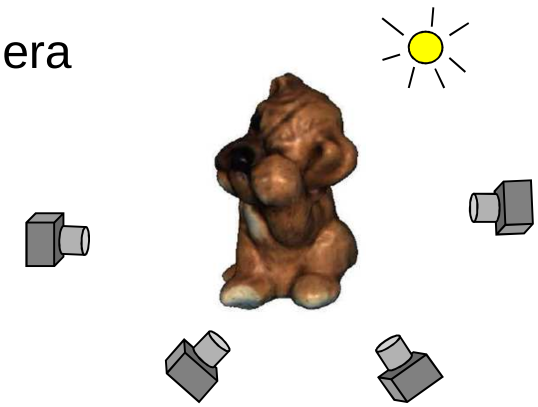

Move object

PS

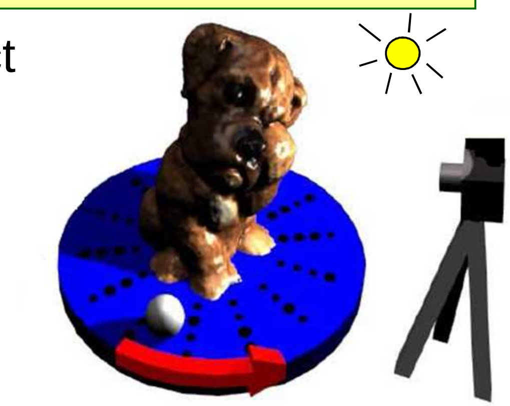

Surface representation

Image-centered

- Depth/disparity w.r.t. to image plane

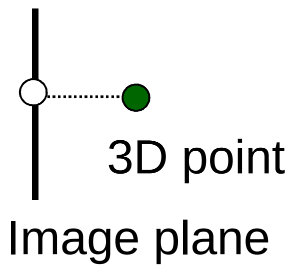

Partial object reconstr.

Limited resolution

Viewpoint dependent

Object-centered

- Voxels

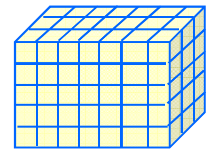

- Level sets (implicit)

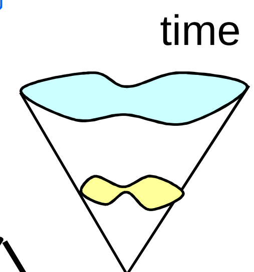

- Mesh

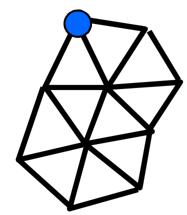

- Depth with respect to a base mesh

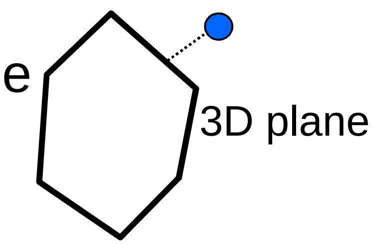

- Local patches

Comparison of different methods

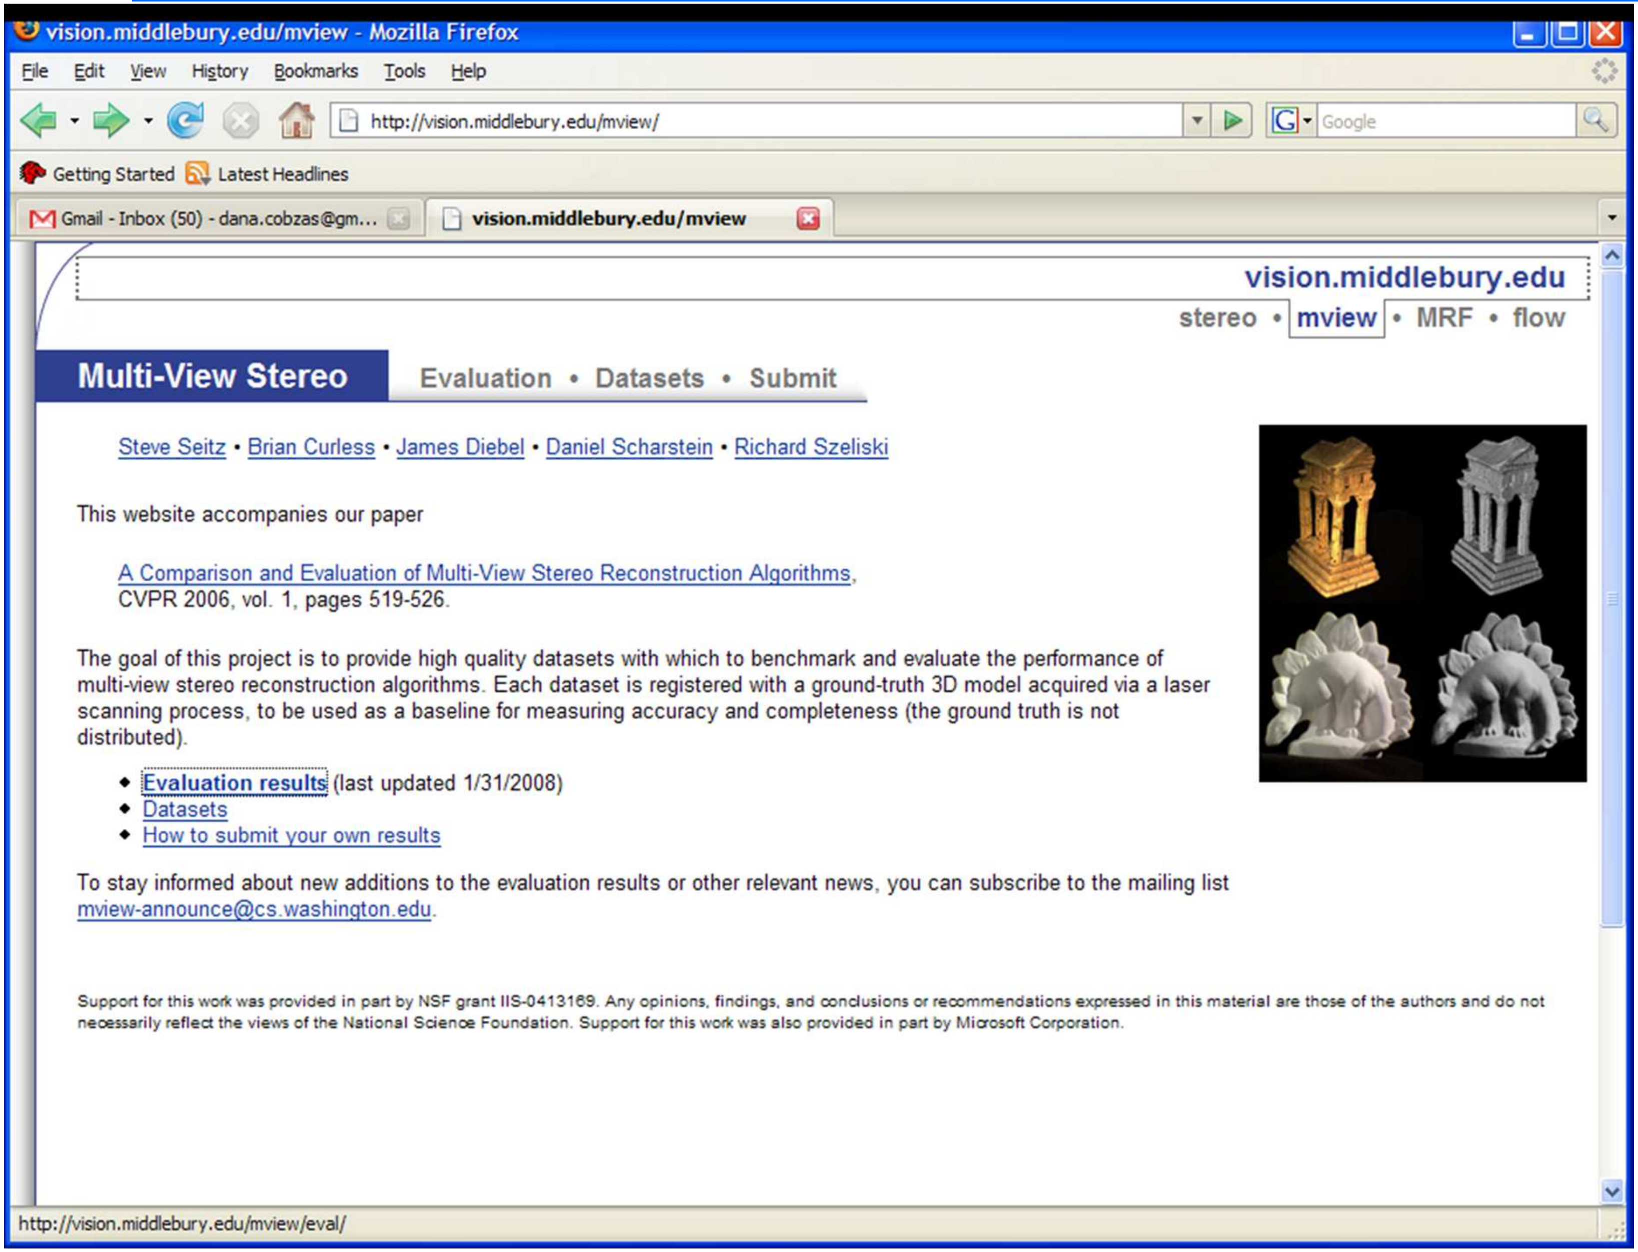

2 datasets

No light

(moving camera)

Volumetric representation

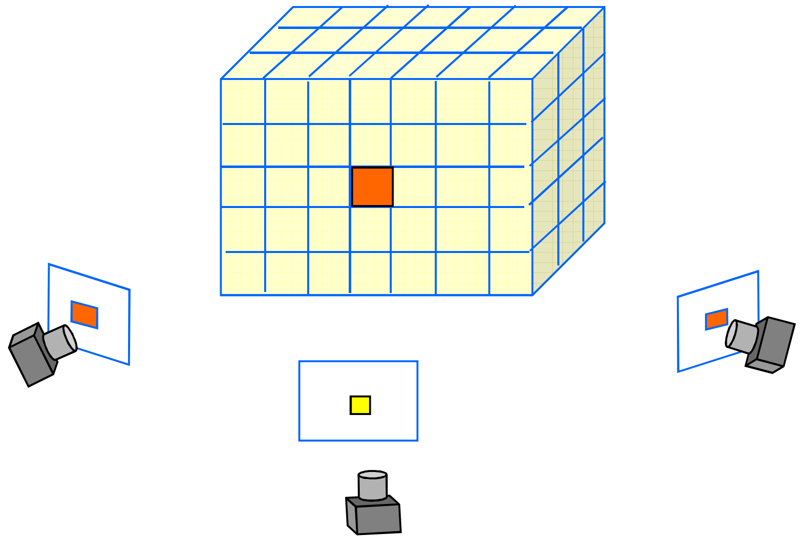

Object : collection of voxels
Normals : ?
Method: carve away voxels that are not photo-consistent with images (purely discrete)
Regularization: no way of ensuring smoothness
Visibility: ensured by the order of traversal (in general a problem!)

Disparity/Depth map

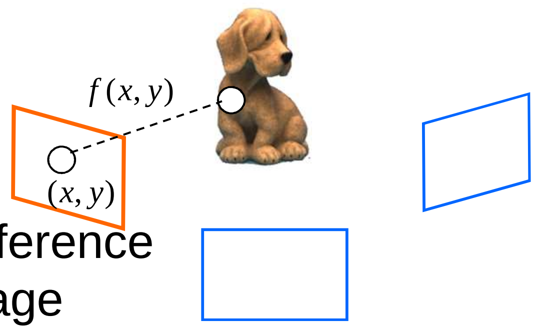

Reference image

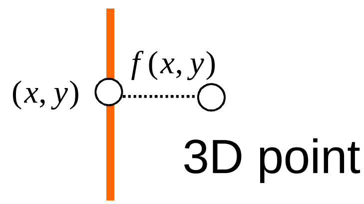

Reference image plane

Object (surface) : $$s(x,y) = (x, y, f(x,y))$$

Normals : $$\mathbf{n} = \frac{\partial s}{\partial x} \times \frac{\partial s}{\partial y} = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, -1 \right)^T$$

Method: find f that best agrees with the input images (minimize the cost functional integrated over the surface)

$$\min_f \iint_{xy} \phi(x, y, f, \nabla f) dxdy$$

Regularization: smoothness on ∇f $\qquad \phi_{smooth} = \nabla f(x,y)^2$

Visibility: ? (mesh defined on image plane + Zbuffering)

Depth w.r.t. base mesh

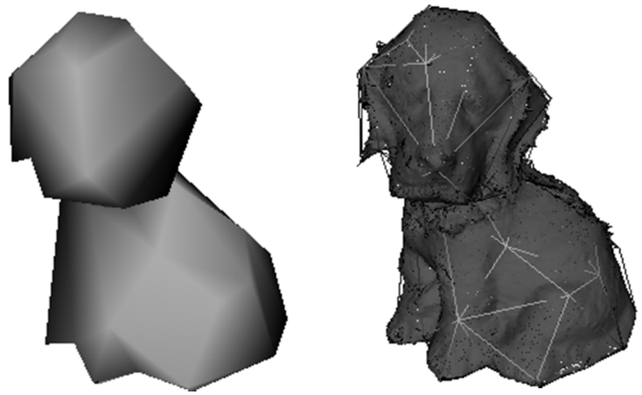

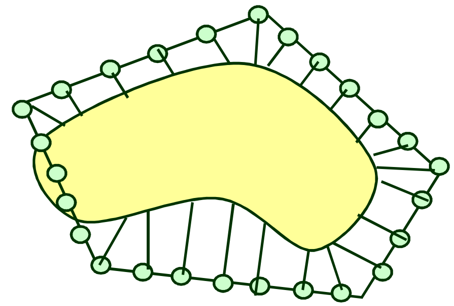

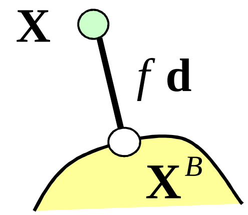

Object (surface) : $\mathbf{X} = \mathbf{X}^B + f\ \mathbf{d}$, d – displacement direction (displacement map)

Normals : local (per triangle) $\mathbf{n} = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, -1 \right)^T$ transform to global CS

Method: $\min_f \int_{X^B} \phi(X^B, f, \nabla f)\, dxdy$

Regularization: smoothness on ∇f (local / global)

Visibility: ? (fine mesh to connect points on each plane + Zbuffering)

How to deal with boundaries ?

Mesh

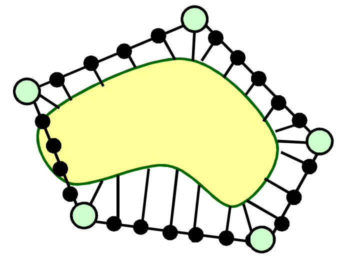

Object (surface) : mesh vertices

$$\mathbf{X} = \lambda_1 \boxed{\mathbf{v}_1} + \lambda_2 \boxed{\mathbf{v}_2} + \lambda_3 \boxed{\mathbf{v}_3}$$

Normals : interpolated

$$\mathbf{n} = \lambda_1 \mathbf{n}_1 + \lambda_2 \mathbf{n}_2 + \lambda_3 \mathbf{n}_3$$

$$\mathbf{n}_j = \sum_{j,k,i \in \Delta(\mathbf{v}_j)} (\mathbf{v}_k - \mathbf{v}_j) \times (\mathbf{v}_i - \mathbf{v}_j)$$

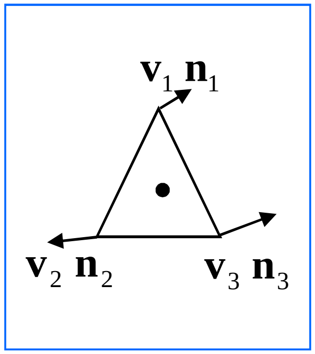

Method: move vertices along interpolated normals based on photo-consistency of neighboring triangles.

Regularization: smooth normals

Visibility: Zbuffering

Topologiocal changes ?

Mixed Representations:Local patches

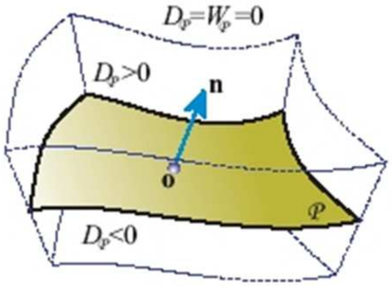

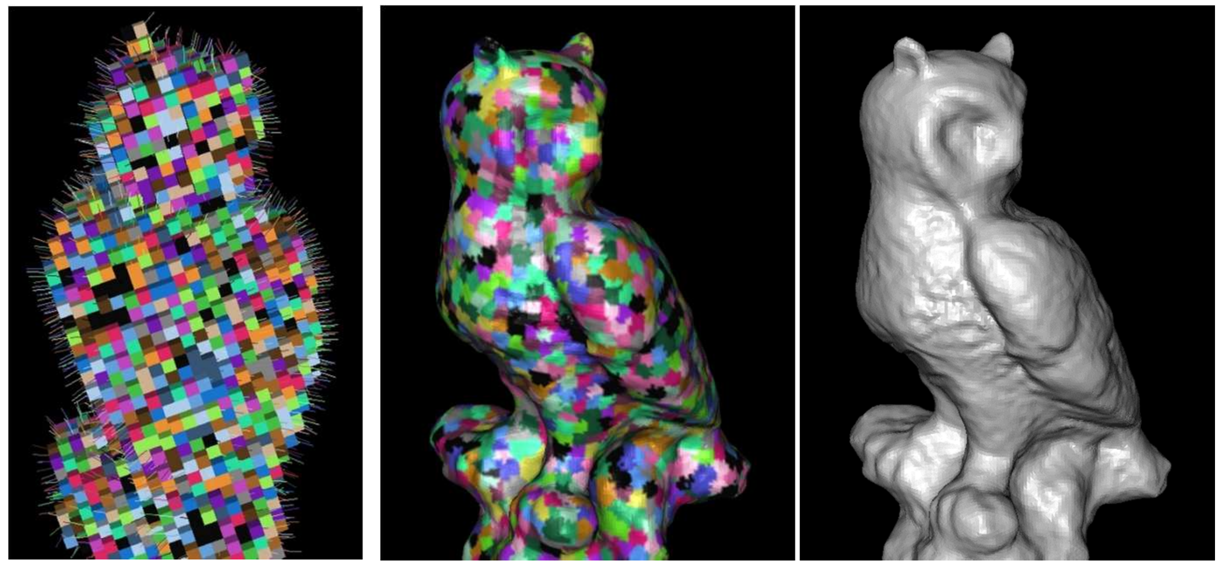

Mixed approaches :

- patches on voxels for a finer surface representation
- mesh on pre-computed voxel correlation (potential fields)
  [Esteban and Schmitt CVIU 2005]
- (depth on base mesh)

Patches :

- arbitrary : [Zheng, Paris et al. IJCV2006]
- quadratic : surfels [Carceroni and Kutulakos IJCV 2002]

Optimization

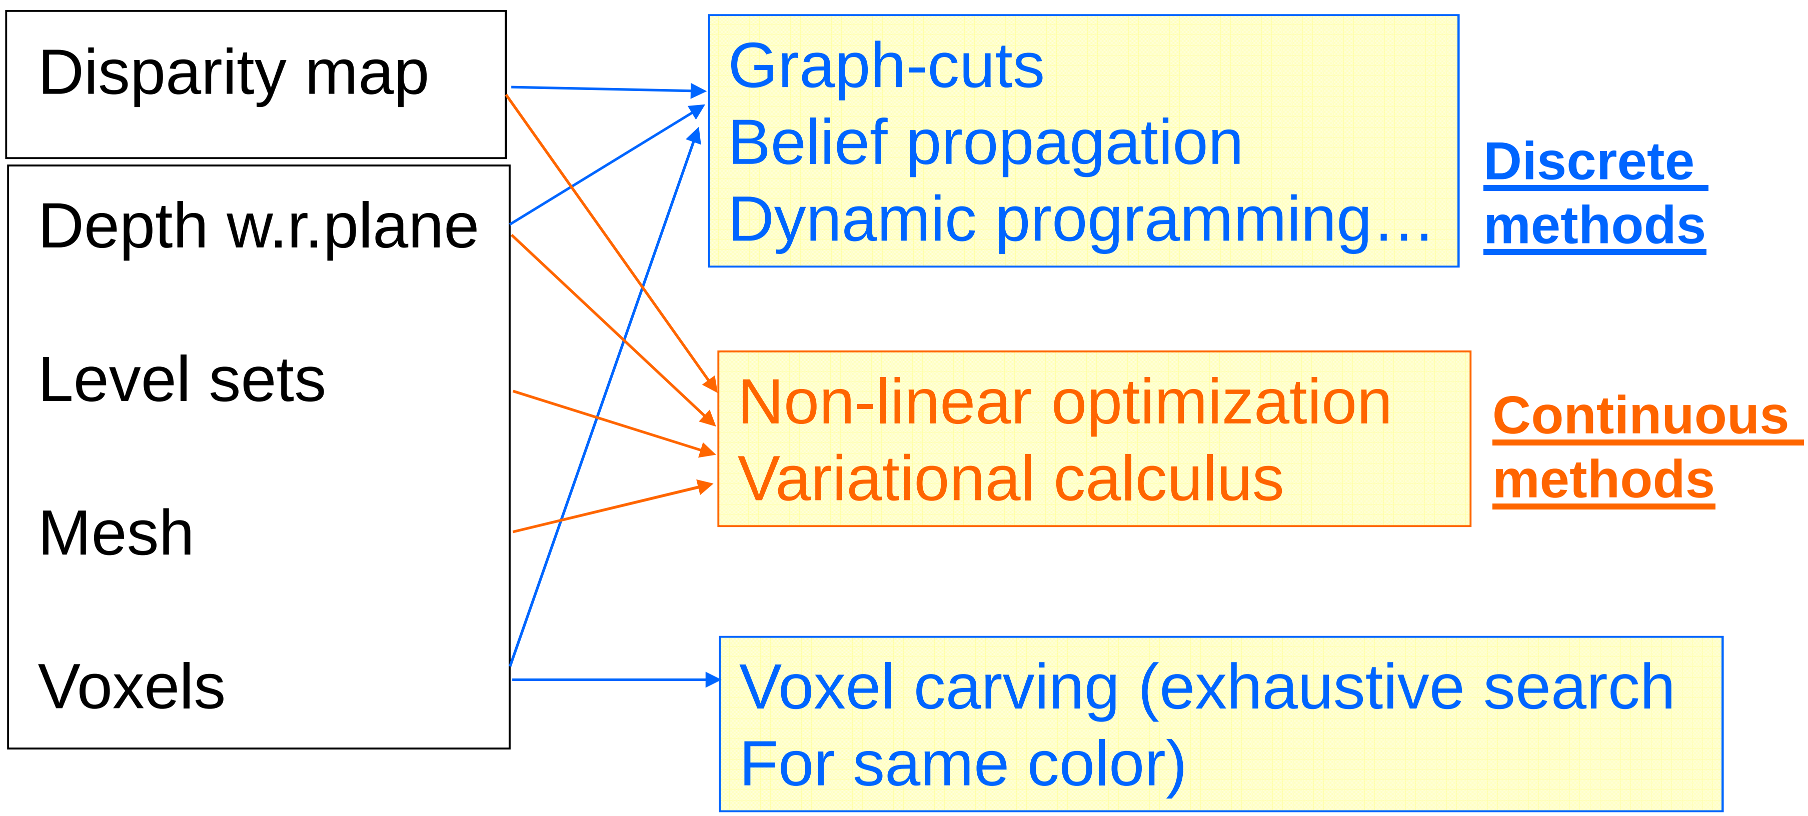

Summary

Discrete

- Voxel carving

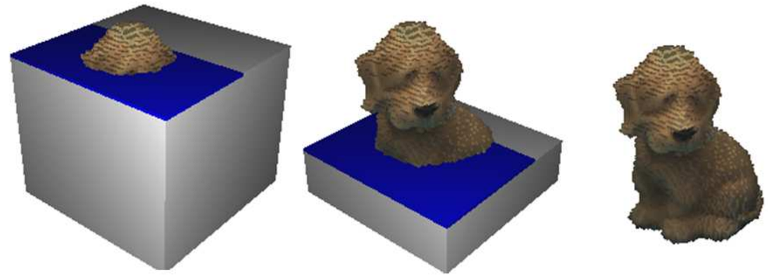

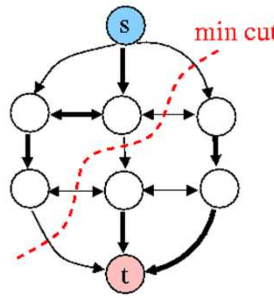

- Graph cut techniques

Continuous

- Variational and level set techniques

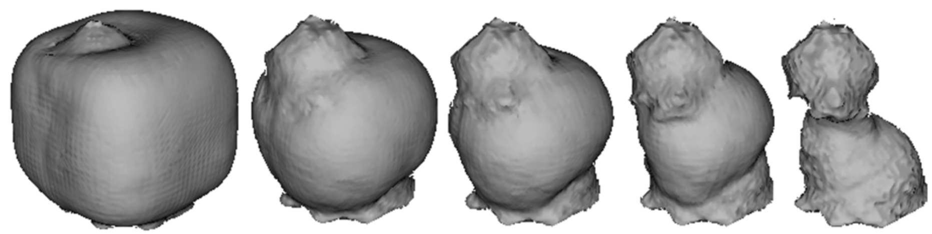

- Mesh-based reconstruction

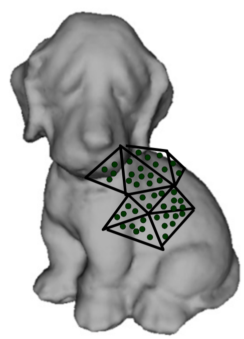

Shape from Silhouette

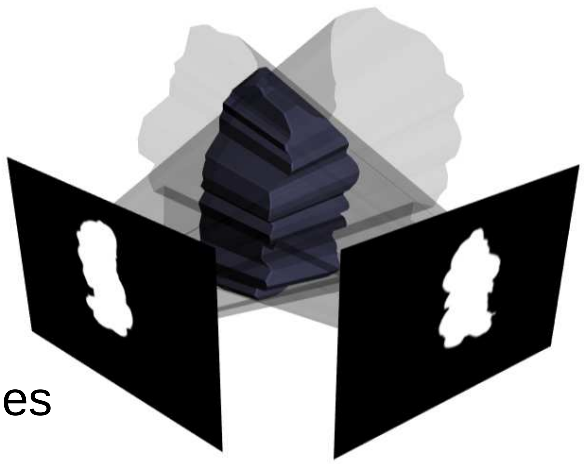

Binary images

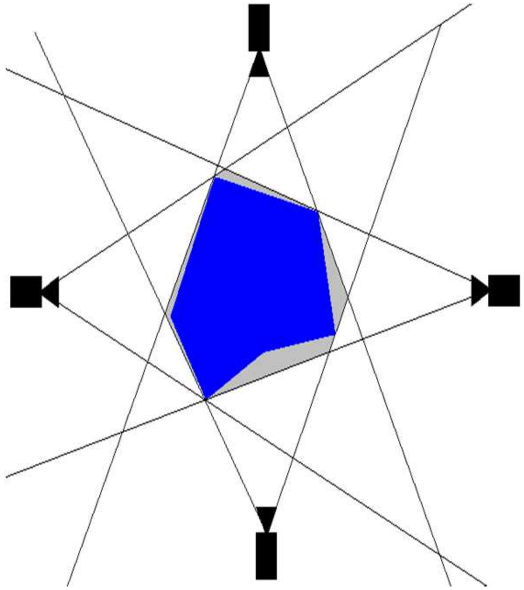

Back project each silhouette $\rightarrow$ 3D cone.

Carve all voxels outside the cone.

Result:

- Reconstruction contains the true scene, but not the same (no concavities)
- Not photo-consistent (only to binary images).
- In the limit reconstructs a **visual hull**

[Martin PAMI 91][Szeliski 93] + Our capgui in lab

Smarter volumetric representations

Voxel-based	Image ray based	Axis-aligned
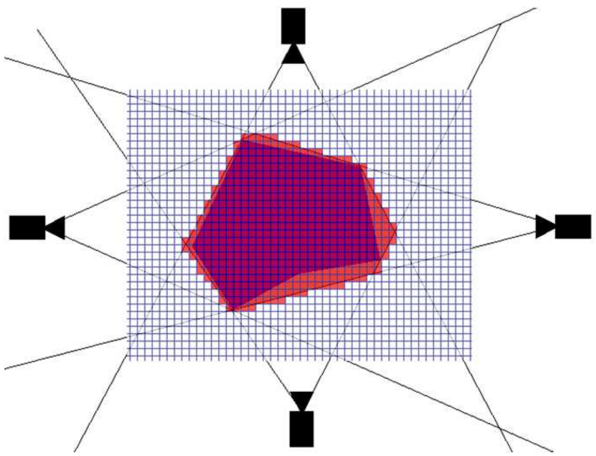	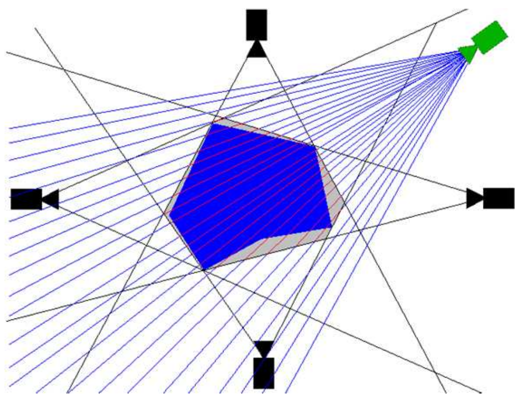	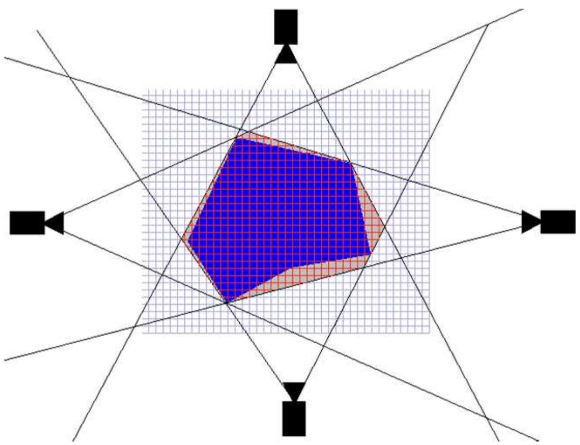
- Inaccurate + Triangulate w. marching cubes	+ Accurate Franco, Boyer	+ Moderatly accurate + Fast + Marching intersections Tarini’01 + our Capgui

Volumetric reconstruction

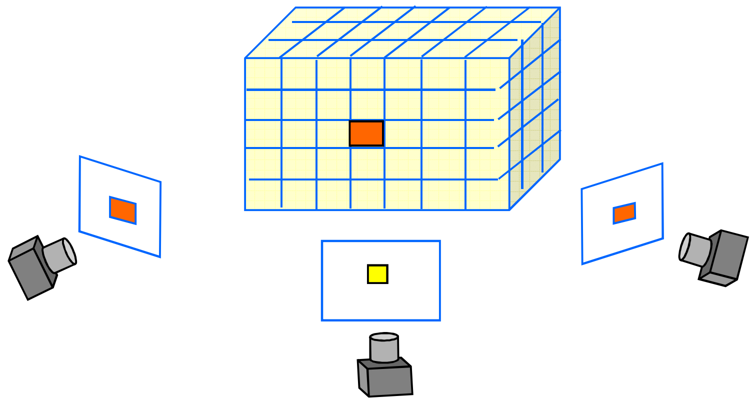

Method : Carve voxels that are not consistent with images (according to a chosen photo-consistency score).

OR

Assign colors to voxels consistent with the input images.

(color + opacity)

N^3 voxels
C colors

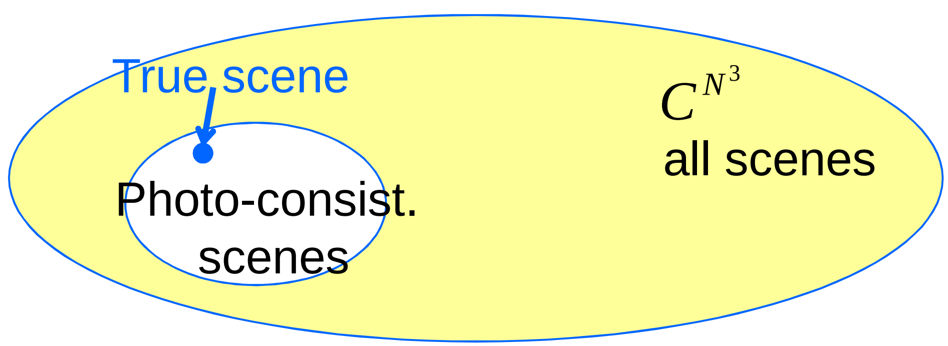

Not unique solution

Order of traversal v. important

Voxel coloring

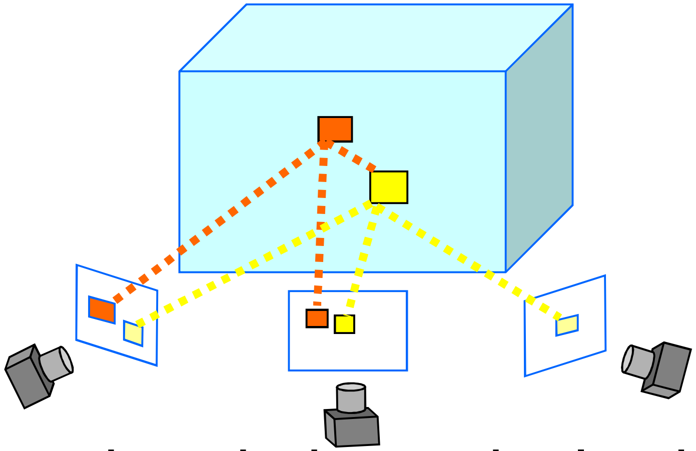

Each voxel

Project in images and correlate

Color if consistent

Visibility !

Depth ordering : Single visibility ordering for each view

(Restriction in camera placement)

Plane sweep [Collins CVPR 96]

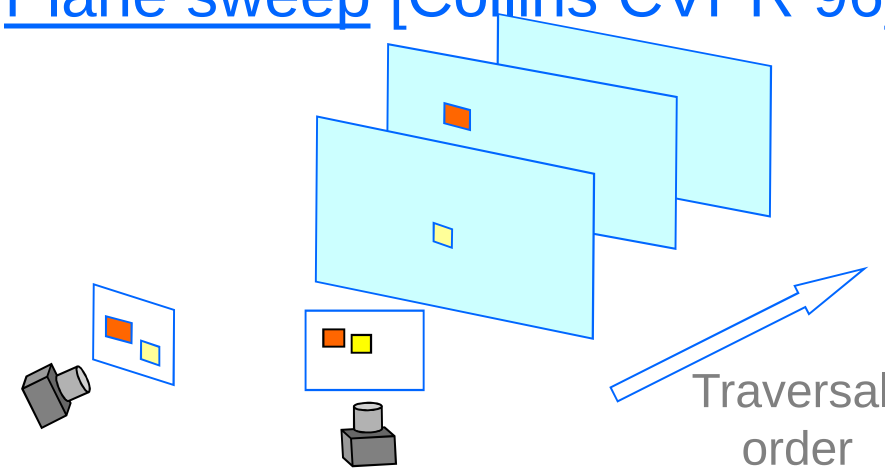

[Seitz,Dyer CVPR97]

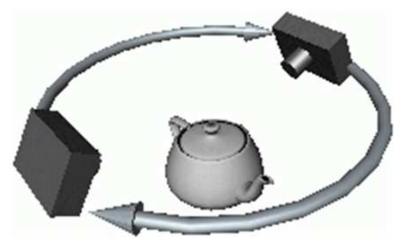

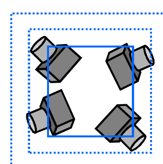

No scene point contained within the convex hull of the cameras

Voxel coloring : results

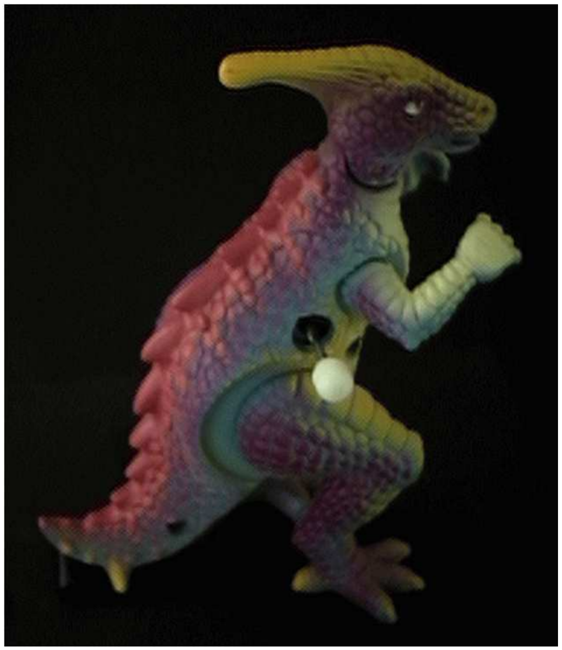

Input image

Results

[Seitz,Dyer CVPR97]

Space carving

In general a view independent order might not exist

Space carving **[Kutulakos, Seitz ICCV 99, IJCV 2002]**

- initialize a volume containing the scene
- Choose a voxel on the surface of the scene
- Project in all visible images
- Carve if not consistent
- Repeat until convergence

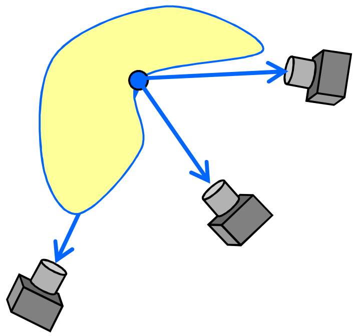

Consistency:

The resulting shape is photo-consistent
(*all inconsistent voxels are removed*)

Convergence:

Carving converges to a non-empty shape
(*a point on the true surface is never removed*)

Photo-hull

Space carving : photo hull

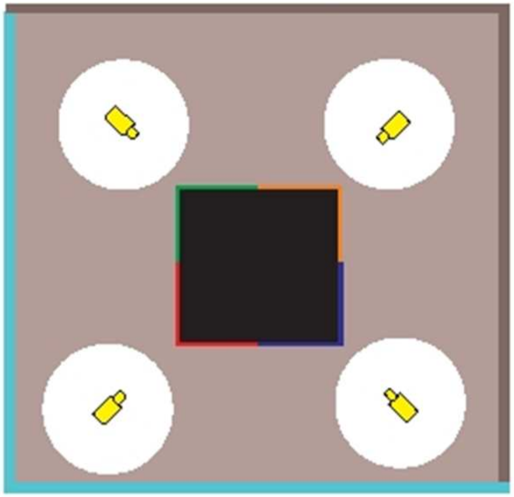

Initial volume and
true scene

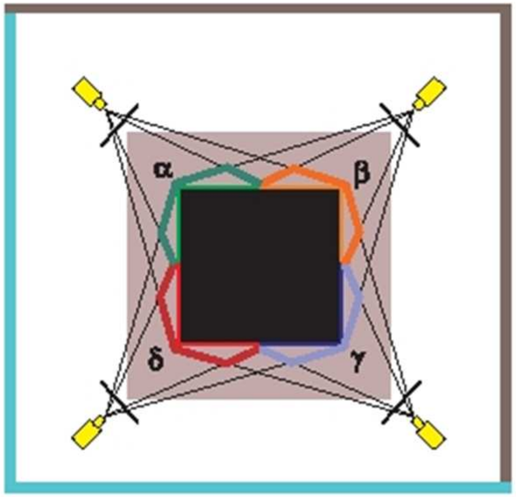

Photo-hull

Photo-hull = union of all photo-consistent shapes

Basic algorithm : requires a difficult update procedure
(visibility computation after carving a voxel)

Multi-pass plane sweep :

- Sweep plane in each 6 directions
- Consider active only cameras on one side of the plane

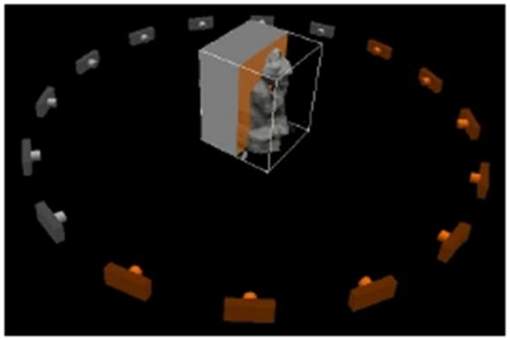

Space carving : results

[Kutulakos, Seitz ICCV 99, IJCV 2002]

Graph cuts for multi-view reconstruction

- Discrete surface reconstruction
- Graph cut
- Graph cuts as hypersurfaces
  - Example: [Paris, Sillion, Quan IJCV 05]
- Types of energies minimized with graph cut

  [Kolmogorov Zabih ECCV 2002]
- Graph cuts for multi-labeling

Reconstruction as labeling

$$\min_S \left(\underbrace{\int_S \phi_{data}(X)dX}_{\text{Photo-consistency}} + \int_S \underbrace{\int_{N(X)} \phi_{smooth}(X,Y)dYdX}_{\text{Smoothness}} \right)$$

Discrete formulation:

- surface representation $p \in P$
- labels f_p

Ex: disparity map — pixels, disparities, $\{0,1,2\ldots\}$

voxels — voxels, occupancy, $\{0,1\}$

Find a set of labels $f = (f_1, \ldots, f_p, \ldots, f_{|P|})$ that minimize

$$E(f) = \sum_p \phi_{data}(f_p) + \lambda \sum_{\{p,q\} \in N} \phi_{smooth}(f_p, f_q)$$

Notes:

- NP hard
- Can be solved using MRF energy minimization methods

graph cuts (submodular E), dynamic programming, belief propagation, simulated annealing …

Graph cuts

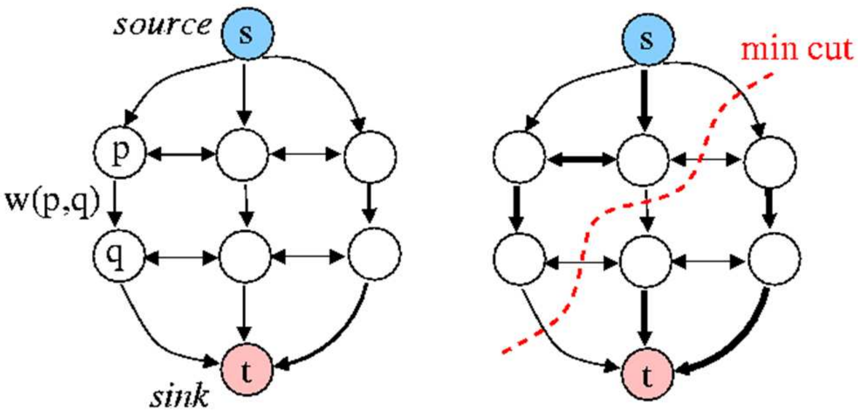

Oriented graph $G = \langle V, E \rangle$

- nodes V, source s, sink t
- edges E, edge capacity $w(p,q)$ (non-negative)

Cut $C=\{S,T\}$ partition of nodes into two disjoint sets such that $s \in S$, $t \in T$

Cost of the cut $\sum_{p \in S, q \in T} w(p,q)$

Minimum cut cut that has minimum cost among all cuts (binary labeling)

Maximum flow maximum amount of liquid that can be sent from the source to the sink interpreting edges as pipes with capacity w.

Polynomial time algorithms

Augmenting paths [Ford & Fulkerson, 1962]

Push-relabel [Goldberg-Tarjan, 1986]

Energy minimization via graph cuts

Motivation:

- Geometric interpretation

  cut = hypersurface in N-D space embedding the corresponding graph
  used to compute optimal hypersurface
- Powerful energy minimization tool for a large class of binary and non-binary energies

  global minimum; strong local minimum

Surface reconstruction:

- Chose a surface representation
- Define a graph (nodes, weights) such that the cost of a cut corresponds to the surface energy function.

How to find global labeling using graph cut ?
What kind of energy can be minimized with a graph cut ?

Graph Cuts as Hypersurfaces

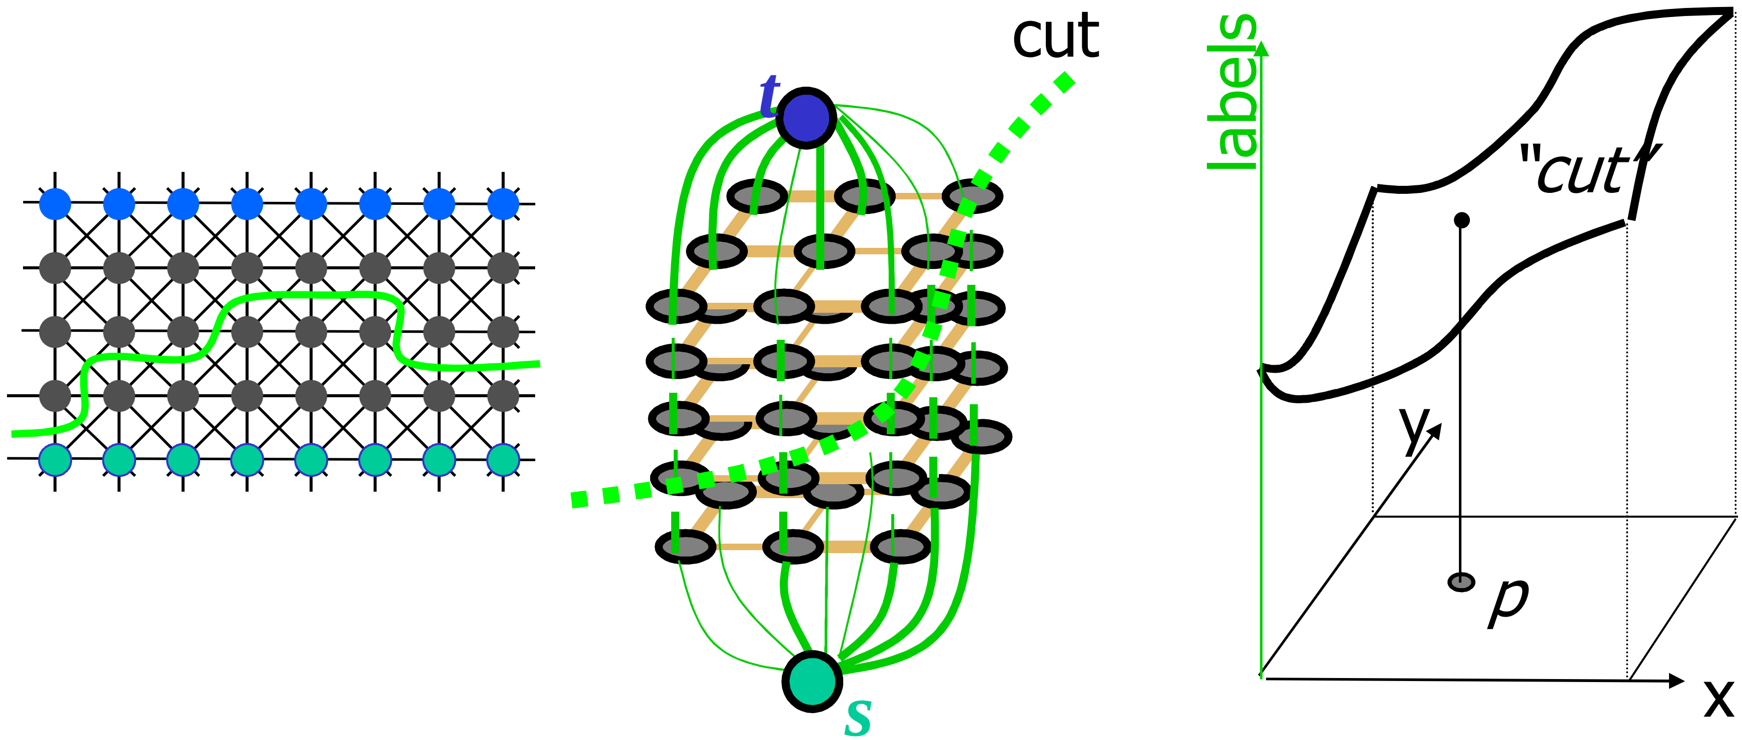

- Graph fully embedded in the working geometric space
- Feasible cut = separated hypersurface in the embedding continuous manifold

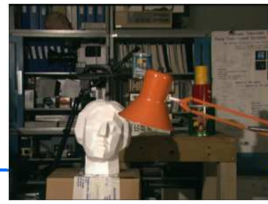

Comparison to “stereo”

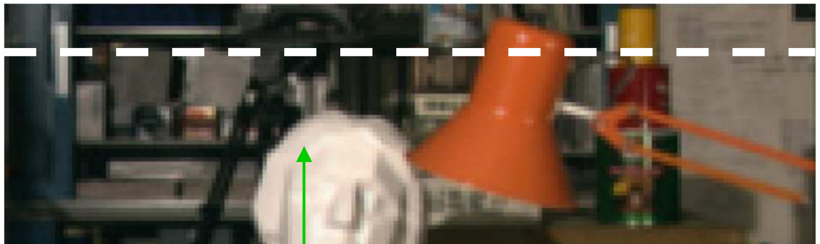

Multi-scanline optimization

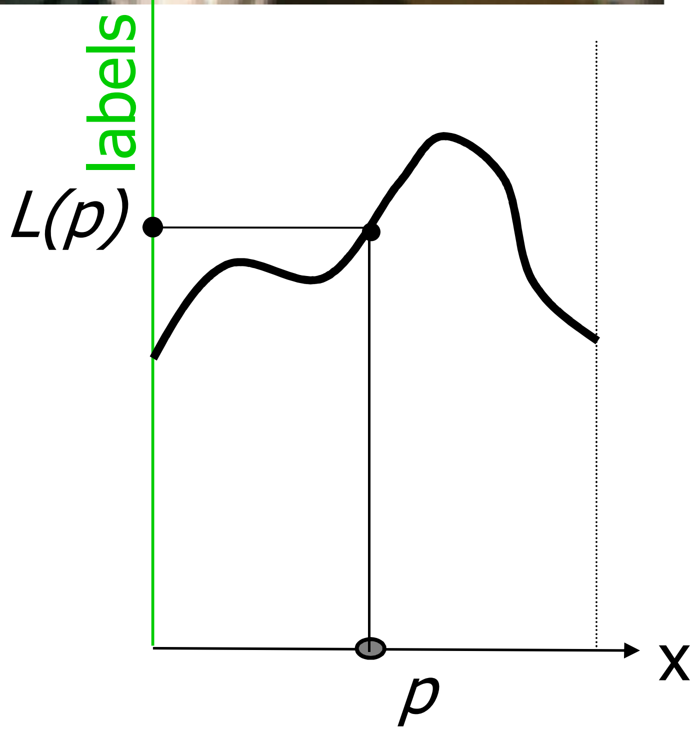

single scan line optimization

eg Search, *Dynamic Programming*

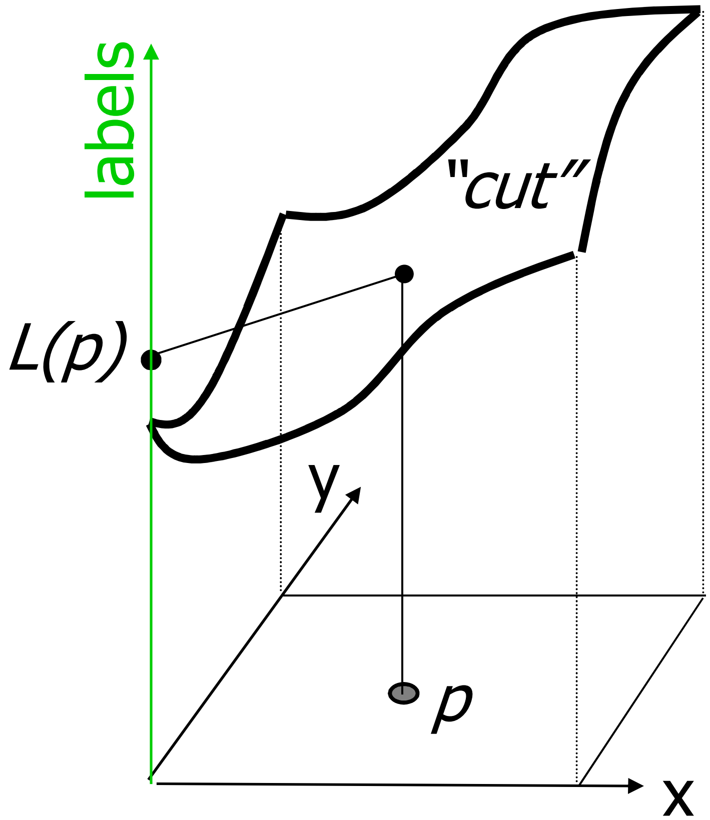

s-t Graph Cuts

multi-scan-line optimization

Example of geometric graph

[Paris, Sillion, Quan: A surface reconstruction method using global graph cut optimization IJCV 05]

Disparity map: pixel $p=(x,y)$ label = disparity $p \rightarrow f_p$

$$E(f)=\sum_p D_p(f_p)+\lambda\sum_{\{p,q\}\in N}\left|f_p-f_q\right|$$

$$D_p(f_p)=I_1(p)-I_2(p+f_p)$$

$$\lambda'=\lambda\left|f_2-f_1\right|$$

$$E_{\min}=D_1(d_4)+\lambda\left|d_4-d_3\right|+D_2(d_3)+D_3(d_3)+\lambda\left|d_4-d_3\right|+D_4(d_4)$$

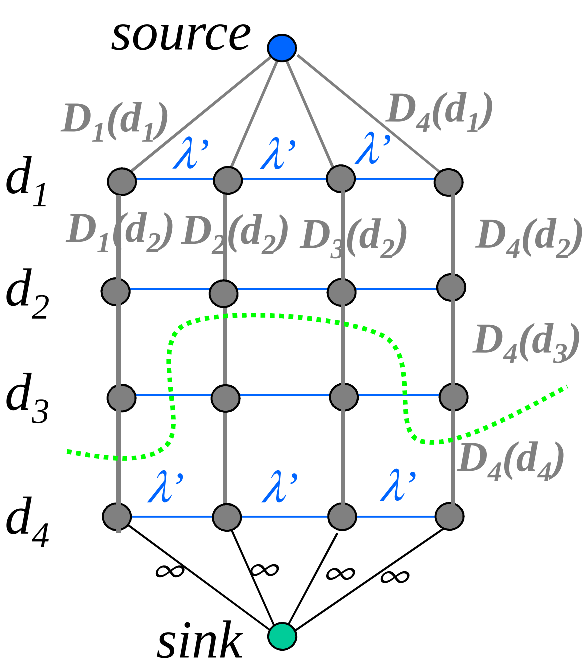

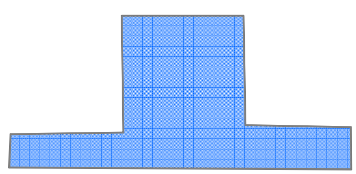

$f_1=d4$
$f_2=d3$
$f_3=d3$
$f_4=d4$

Graph

Surface

Results

Convex smoothing – global solution

[Paris, Sillion, Quan IJCV 05, ACCV 04]

Types of regularization energies

$$E(f) = \sum_p D(f_p) + \lambda \sum_{\{p,q\} \in N} V(f_p, f_q)$$

Convex

+ global convergence

- oversmooth

$$V(f_p, f_q) = |f_p - f_q|$$

$$V(f_p, f_q) = (f_p - f_q)^2$$

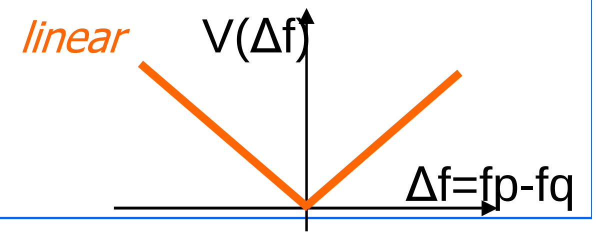

Preserves discontinuities

NP hard

? convergence

$$V(f_p, f_q) = \min(T, |f_p - f_q|)$$

$$V(f_p, f_q) = \min(T, (f_p - f_q)^2)$$

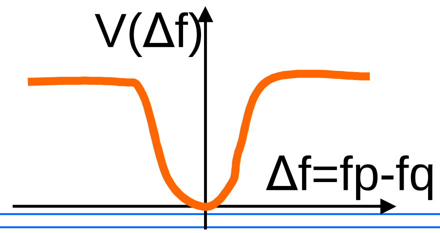

Potts

piecewise planar

binary graph

$$V(f_p, f_q) = \lambda_{pq} \rho[f_p \neq f_q]$$

$$\rho(b) = \begin{cases} 1, true \\ 0, \ false \end{cases}$$

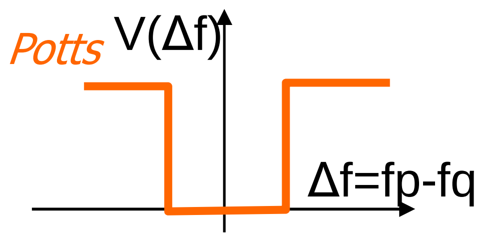

Graph cuts : state of the art

[Boykov CVPR 05 Tutorial]

- Optimization of first-order properties of segmentation boundary
  (Riemannian length/area, flux of a vector field)
  Can’t optimize curvature of the boundary (for now)

- Class of energies that can be **minimized exactly**
  - **binary** energies with regular (sub-modular) interactions
  - **multi-label** (non-binary) energies with “convex” interactions
    - excludes robust *discontinuity-preserving* interactions

- Guaranteed quality **approximation** algorithms for **multi-label** energies with **discontinuity-preserving** interactions…
  - Potts model of interactions
  - *Metric* interactions
  - Regular (sub-modular) interactions

GOTO slide 41

Exact multi-labeling

Linear and convex smoothing (interaction) energy
Geometric graphs.

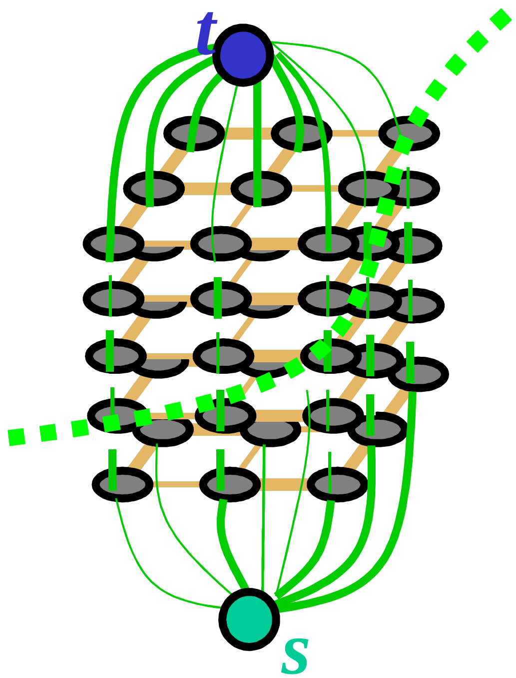

- Multi-scan-line stereo
  - [Roy & Cox 1998, 1999]
  - [Ishikawa & Geiger 1999] (occlusion handling)
- “Linear” interaction energy
  - [Ishikawa & Geiger 1998]
  - [Boykov, Veksler, Zabih 1998]
- Convex interaction energy
  - [Ishikawa 2000, 2003]

Binary graphs

[Kolmogorov, Zabih : What energy functions can be minimized via graph cuts? ECCV 2002]

$$E(f) = \sum_p D_p(f_p) + \sum_{pq \in N} \lambda_{pq} \cdot \delta_{fp \neq fq}$$

$$f_p \in \{s, t\}$$

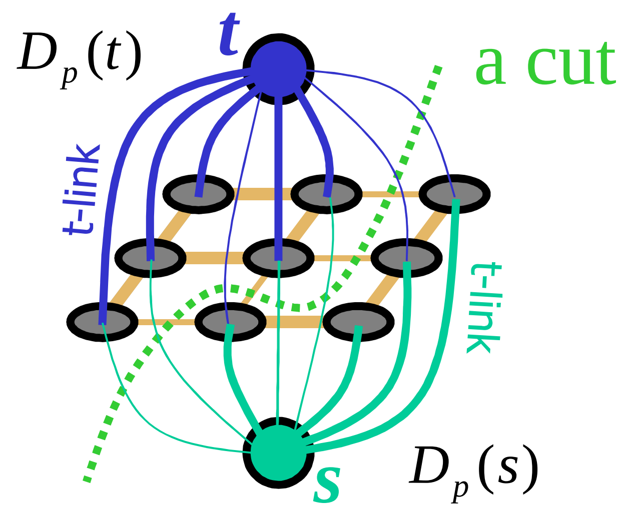

Complete characterization of energies that can be minimized with graph cut :

E(f) can be minimized by *s-t* graph cuts $\Leftrightarrow$ $V(s,s) + V(t,t) \leq V(s,t) + V(t,s)$

Large class of energies (Potts, metric …)

BUT multi-view reconstruction is a multi-labeling problem !

Approximate multi-labeling

<u>Binary Potts energy</u>
Extended to multi-labeling
NP hard (≥ 3 labels)

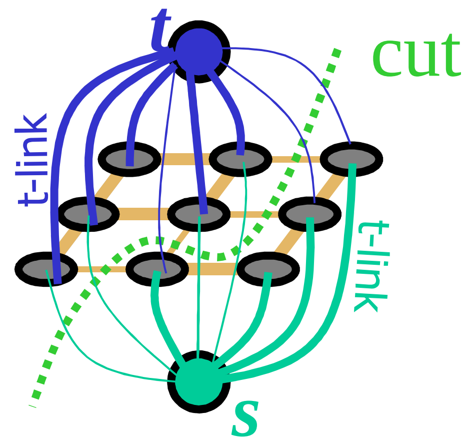

[Boykov et al: Fast approximate energy via graph cuts, PAMI 2001]

α <u>Expansion</u> approximate solution

- Ideea: break optimization into a set of binary s-t cut problems
- Each iteration – consider one label *a*
- Binary cut: some labels are relabeled with *a*; the others remain unchanged

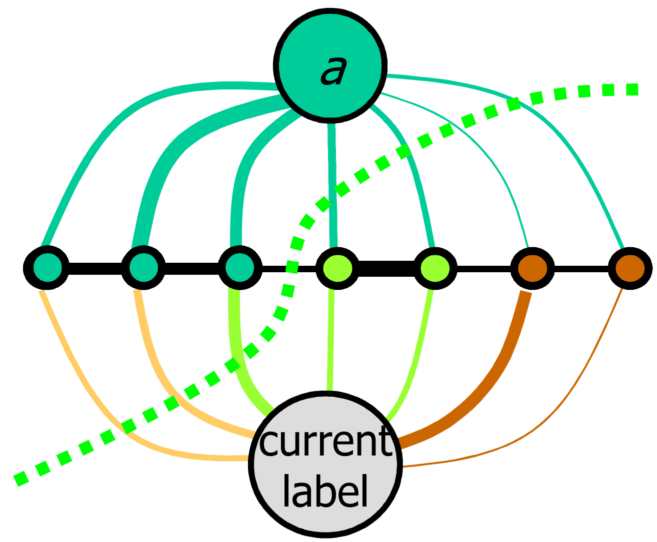

α - expansion

Examples of standard and large moves from a given iniotia; labeling.

The number of labels is |L| = 3.

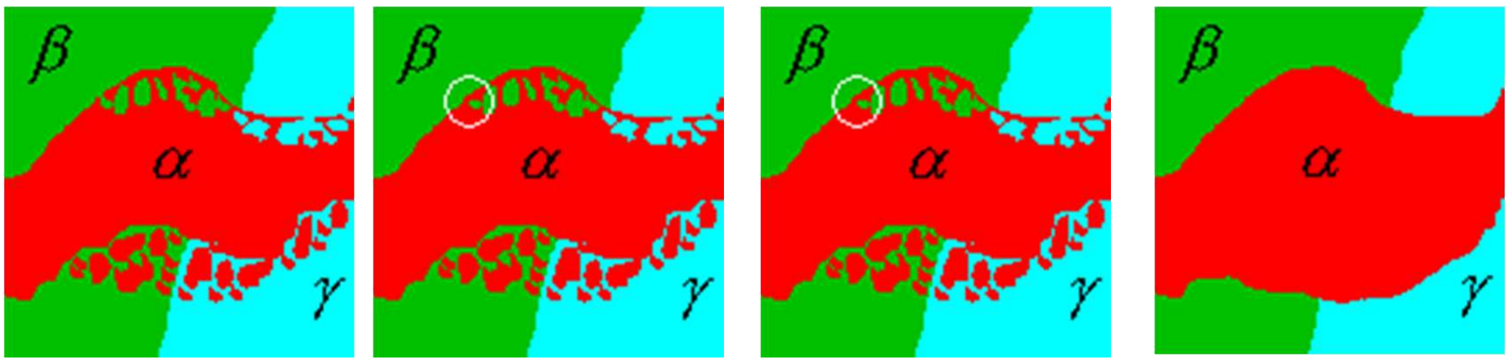

initial labeling standard move swap a-expansion

α expansion

Properties

Guaranteed approximation quality

- within a factor of 2 from the global minima (Potts model)

Applies to a wide class of energies with robust interactions

- Potts model
- "Metric" interactions
- "Submodular" (regular) interactions

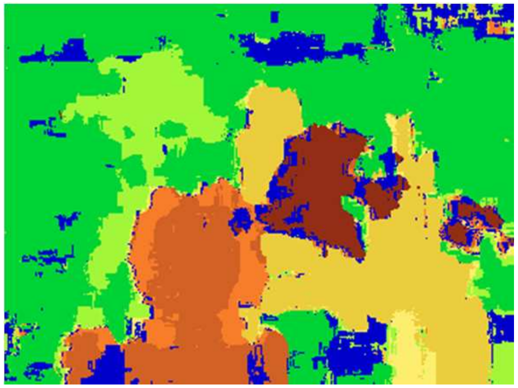

normalized correlation,
start for annealing, 24.7% err

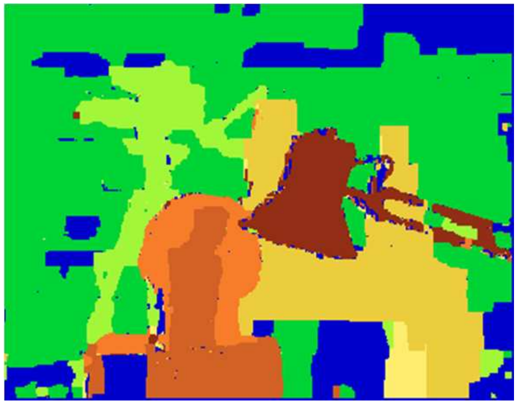

simulated annealing,
19 hours, 20.3% err

a-expansions (BVZ 89,01)
90 seconds, 5.8% err

Graph cut Example

[Vogiatzis, Hermandez-Esteban, Torr, Cipolla PAMI 2007, CVPR 2005]

Surface representation: voxels; no need for bounding inner/outer surface
Regularization : weighted volume – balloon force
Minimization : graph cut
Occlusions : accounted using a voting photo-consistency score
(occluded pixels are treated as outliers)

$$E[S] = \underbrace{\iint_S \rho(\mathbf{x})dA}_{\text{Photoconsistency}} + \underbrace{\iiint_{V(S)} \sigma(\mathbf{x})dV}_{\text{Foreground/background cost}}$$

S = surface
V(S) = foreground

$\sigma(x)=-\lambda$ balloon force

silhouette cue – make $\sigma(x)$ very large outside VH

Photo-consistency metric

Account occlusions   ρ(x,S) [continuous, level set formulations]
S determines visibility but S is the solution!

Problem : Not suitable for graph-cut

Solution : ρ(x) that accounts for occlusions using NCC

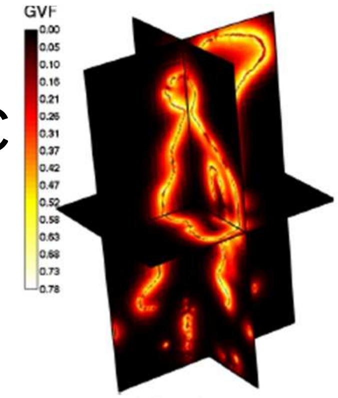

8 levels

$$\rho(\mathbf{x}) = \exp\{-\mu \sum_{i=1}^{N} \mathrm{VOTE}_i(\mathbf{x})\}.$$

Optic ray $\mathbf{o}_i(d) = \mathbf{x} + (\mathbf{c}_i - \mathbf{x})d$

Correlation scores $\mathcal{C}(d) = \sum_{j \in \mathcal{N}(i)} S_j(d).$

Vote $\mathrm{VOTE}_i = \begin{cases} \mathcal{C}(0) & if \quad \mathcal{C}(0) \geq \mathcal{C}(d)\ \forall d \\ 0 & otherwise \end{cases}.$

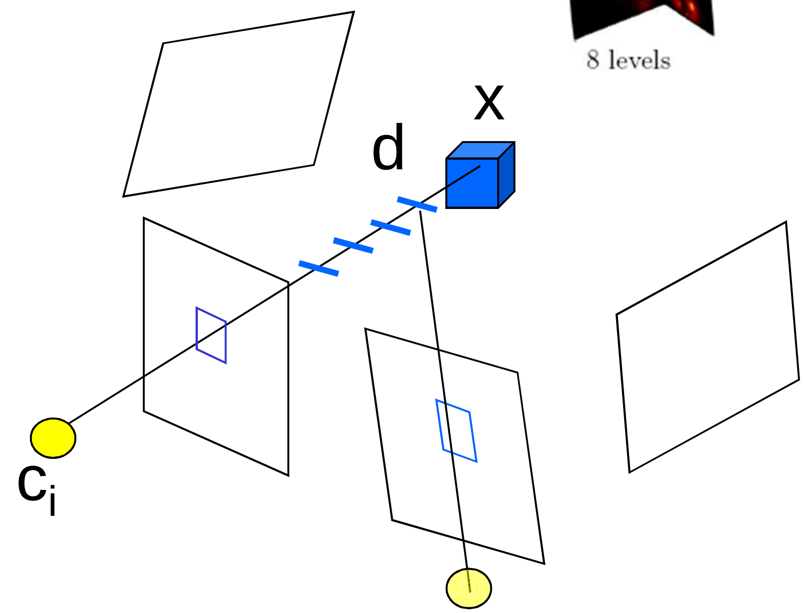

Account occlusions

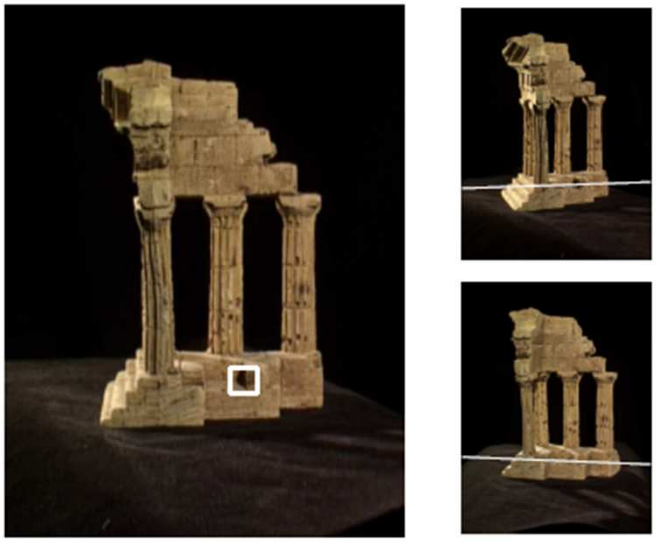

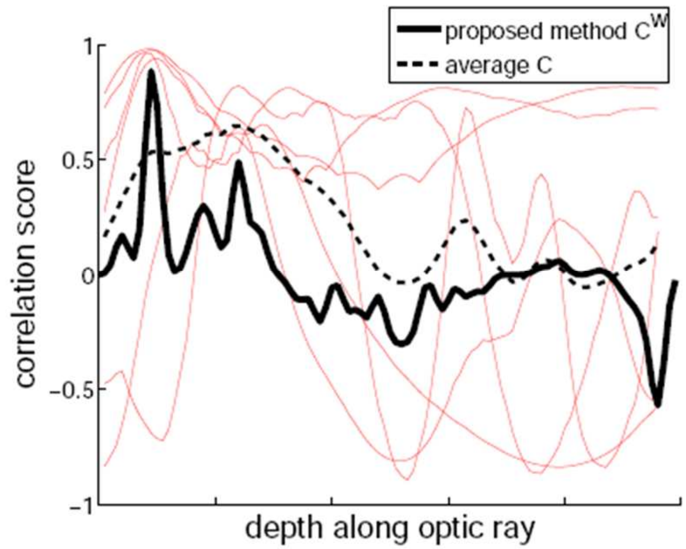

$$\mathcal{C}(d) = \sum_{j \in \mathcal{N}(i)} S_j(d).$$

$$\mathcal{C}^w(d) = \sum_{j \in \mathcal{N}(i)} \sum_k S_j(d_k) W(d - d_k).$$

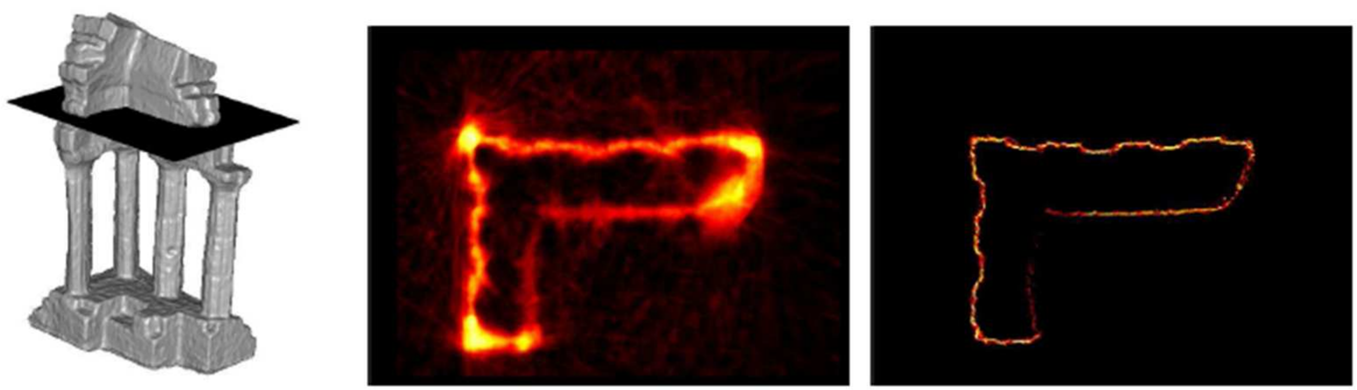

Graph structure

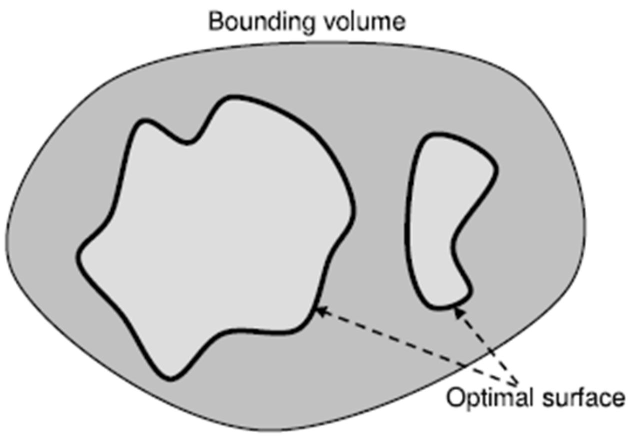

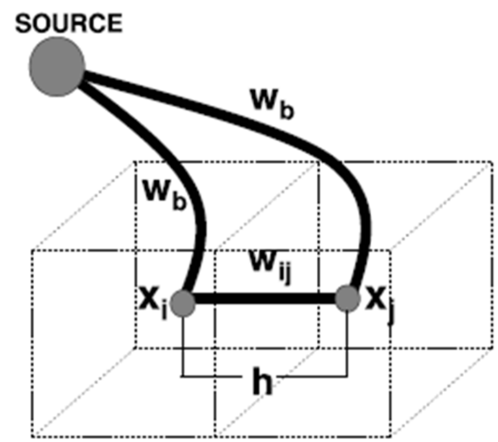

<u>Data (photo-consistency)</u>

$$w_{ij} = \frac{4\pi h^2}{3} \rho \left(\frac{\mathrm{x_i} + \mathrm{x_j}}{2} \right)$$

<u>Ballooning</u>

$$w_b = \lambda h^3$$

6 neighboring system

Results

Vogiatzis2 0.50mm
Furukawa2 0.54mm

Competitor – Yasu Furukawa ☺

Video

Multi-view stereo

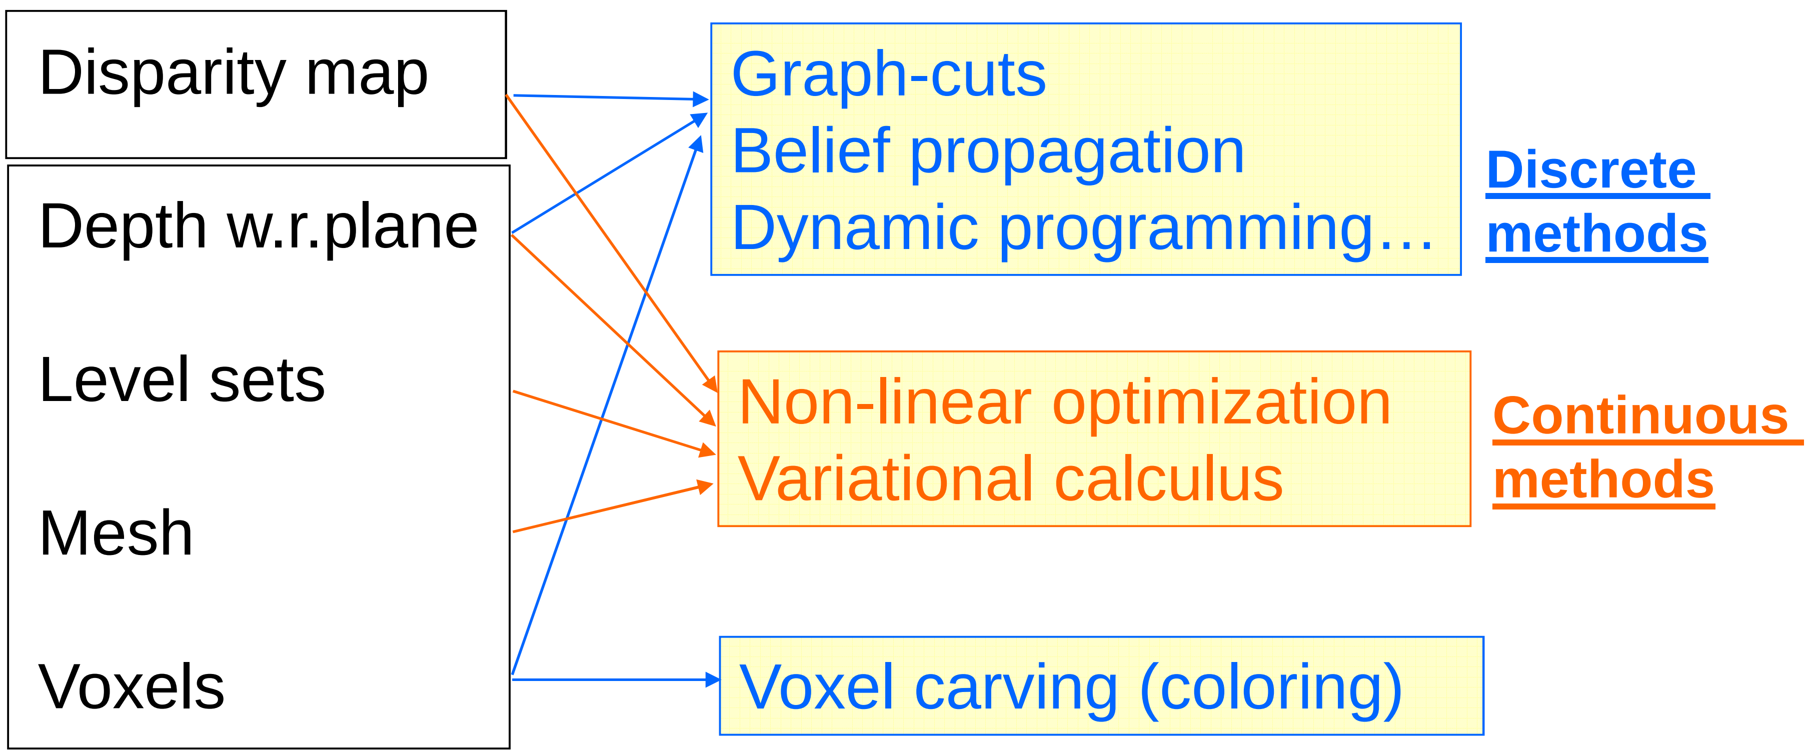

Continuous multi-view methods

- Regular surface and surface evolution
- Level set methods
- Example of mesh-based reconstruction

Surface evolution

Continuous formulation

recover shape (surface) by minimizing cost functional integrated over the surface.

$$\min_S \int_{X \in S} \phi(X)\, dX$$

- Cost functional : photo-consistency + regularization (smoothness)
- Numerical methods: gradient descent, conjugate gradient, level sets …
- Natural extension of curve evolution (2D) [Caselles ICCV95] to 3D [Robert,Deriche ECCV 96][Faugeras Kerivan 98]

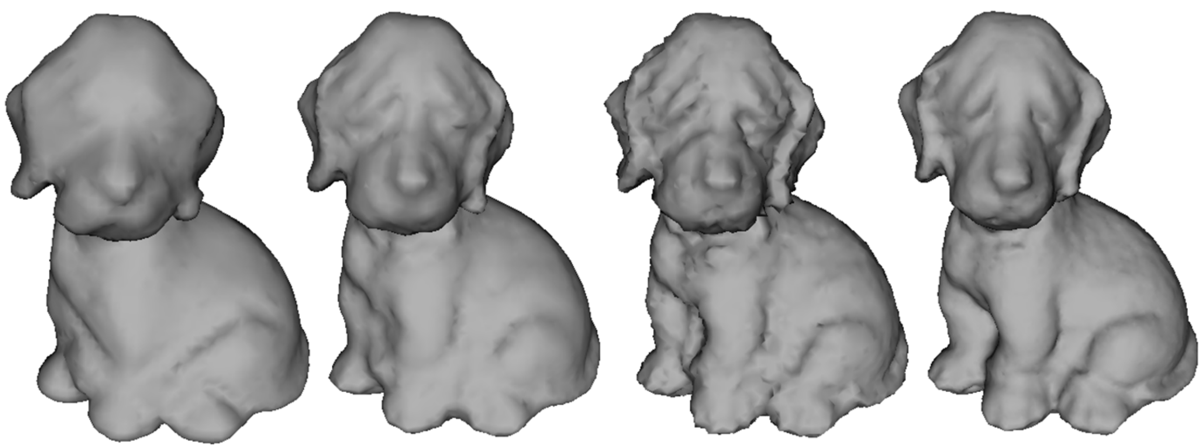

Regular surface = smooth

Definition: Regular surface (manifold)

$S \subset \Re^3$ is a regular surface if for each point $\mathbf{P} \in S$ there exist a neighborhood V and a map $\mathbf{X}: U \to V \cup S$ of an open set $U \subset \Re^2$ such that (**X** = parametrization):

1. $\mathbf{X}$ differentiable

2. $\mathbf{X}$ homeomorphism ($\exists \mathbf{X}^{-1}: V \cup S \to U$ continuous)

3. $\forall (u,v) \in U$ the differentiable $d\mathbf{X}|_{(u,v)}: \Re^2 \to \Re^3$ is one to one

Regular surface - properties

Regular surface S

- tangent vector

$$w \in span\ (\mathbf{X}_u, \mathbf{X}_v) = T_P(S)$$

- normal

$$\mathbf{n}_P = \frac{\mathbf{X}_u \times \mathbf{X}_v}{|\mathbf{X}_u \times \mathbf{X}_v|}$$

- area

$$A(R) = \int_Q |\mathbf{X}_u \times \mathbf{X}_v| dudv = \int_R dA \qquad Q = \mathbf{X}^{-1}(R)$$

- curvature

$$H = \frac{1}{2}\left(\frac{\langle \mathbf{n}, \mathbf{X}_{uu} \rangle}{\langle \mathbf{X}_u, \mathbf{X}_u \rangle} + \frac{\langle \mathbf{n}, \mathbf{X}_{vv} \rangle}{\langle \mathbf{X}_v, \mathbf{X}_v \rangle}\right) \qquad (\mathbf{X}_u \perp \mathbf{X}_v)$$

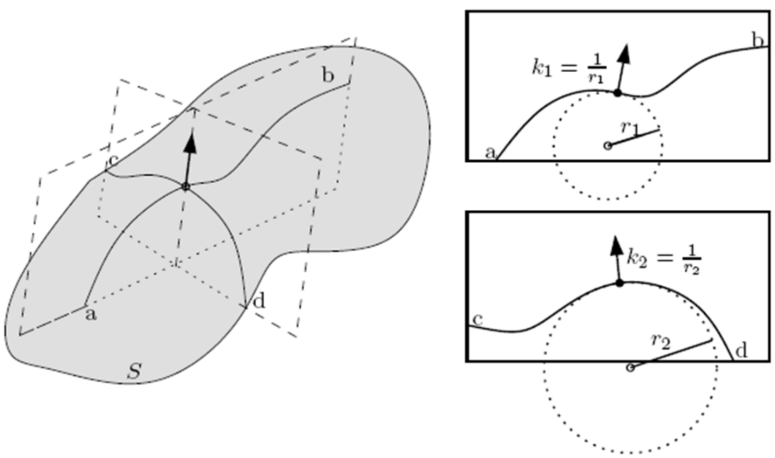

Surface evolution

Cost functional $\Phi : \Re^3 \times \Re^3 \rightarrow \Re^+$

Energy of the surface

$$E = \int_S \Phi(\mathbf{X}, \mathbf{n})\, dA \qquad E = \int_S \Phi(\mathbf{X})\, dA$$

Evolution flow (Euler-Lagrange equations)

$$S_t = \left(2H\Phi - \langle \Phi_X, \mathbf{n} \rangle\right)\mathbf{n}$$

$$S(0) = S_0$$

OBS:

- problem is intrinsic (independent on parameterization) dA
- automatic regularization H
- motion in the normal direction
- whole surface is evolving in time (reference frame attached to the object)

1. Discretization
2. Choice of cost functional

Photo-consistency functional

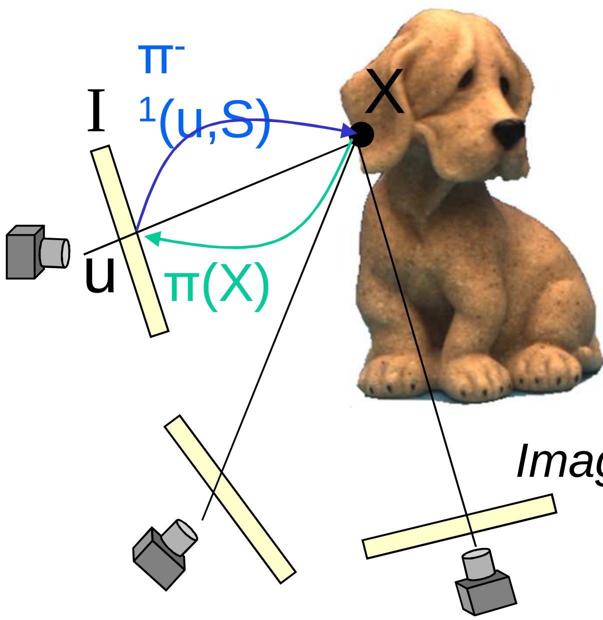

	On Image	On surface
Image point	u	π(X)
Surface point	X(u)=π-1(u,S)	X
	$\Phi_{im}(u) = I(u) - f(X(u))$	$\Phi_S(X) = I(\pi(X)) - f(X)$
Image i:	$E_i = \int_{\boxed{\Omega}} \Phi^i{}_{im}(u)\boxed{du}$	$E_i = \int_{\boxed{S}} \Phi^i{}_S(X)\boxed{dA}$
	$E = \sum_i E_i$	$E = \sum_i E_i$

S-surface,

f(X)-color (radiance) of X

Where to integrate image/surface ?

On Image

$$\Phi_{im}(u) = I(u) - f(X(u))$$

$$E_{im} = \int_{\Omega} \Phi_{im}(u)\,du$$

$$du = \frac{XN}{Z^3}\,dA$$

$$E_{im} = \int_{S} \Phi\ (X) \boxed{\frac{XN}{Z^3}\,dA}$$

$$\frac{dX_{im}}{dt} = \nabla \bullet \left(\Phi \frac{X}{Z^3}\right) N = \left(\nabla\Phi \frac{X}{Z^3} + \cancel{\Phi \frac{\nabla X}{Z^3}}\right) N$$

$$\nabla\Phi = \cancel{\nabla I(I-f)} + \nabla f(I-f)$$

$$\nabla I \bullet X = 0, \quad \nabla X \bullet N = 0$$

$$\frac{dX_{im}}{dt} = \nabla f(I-f) \frac{X}{Z^3} N$$

Constant f $\nabla f = 0$

Impossible to reconstruct

On surface

$$\Phi_S(X) = I(\pi(X)) - f(X)$$

$$E_S = \int_{S} \Phi\ (X) \boxed{dA}$$

$$\frac{dX_S}{dt} = (\underbrace{\nabla\Phi N} + \underbrace{\Phi \bullet H})N$$

1.Depends on image derivative

2.Automatic regularizatior

3.Doesn't account for image discretization

Surface evolution example

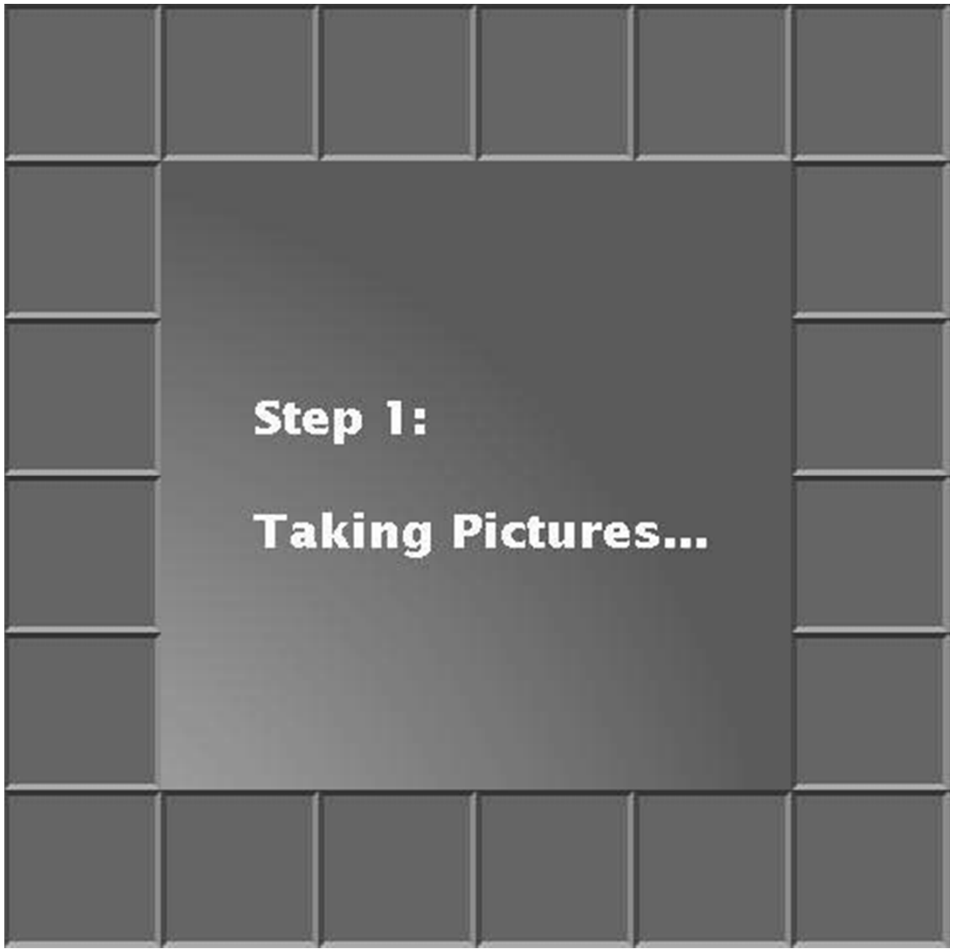

[Faugeras Kerivan 98]

Modeling correct visibility

$$E(\Gamma) = \int_{\mathbb{R}^3} \left(\nabla g(\mathbf{x}) \cdot \frac{\mathbf{x}}{\mathbf{x}_z^3}\right) \nu_\Gamma(\mathbf{x})\, d\mathbf{x}\,.$$

$$\underbrace{(I-C)^t \nabla C \frac{\mathbf{x}}{\mathbf{x}_z^3}\nu_\Gamma}_{\text{Usual term}} + \underbrace{((I-C)^2-(I-C')^2)\frac{\mathbf{x}^t\nabla\mathbf{n}\mathbf{x}}{\mathbf{x}_z^3}\delta(\mathbf{x}\cdot\mathbf{n})\nu_\Gamma}_{\text{Horizon term}}$$

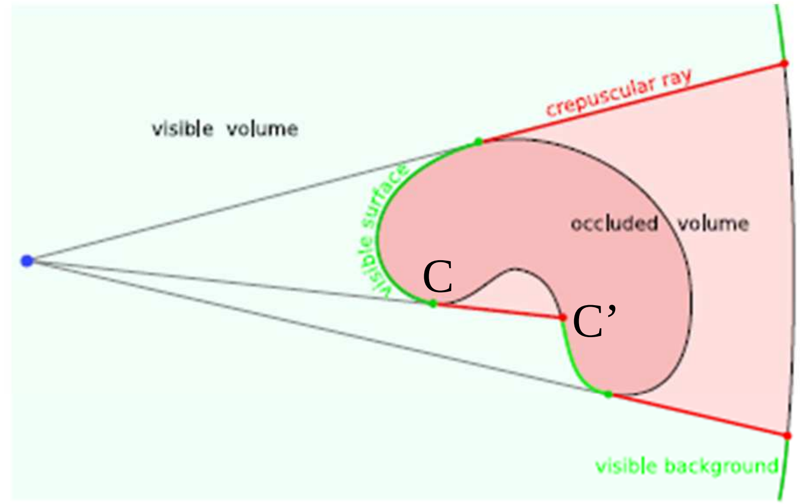

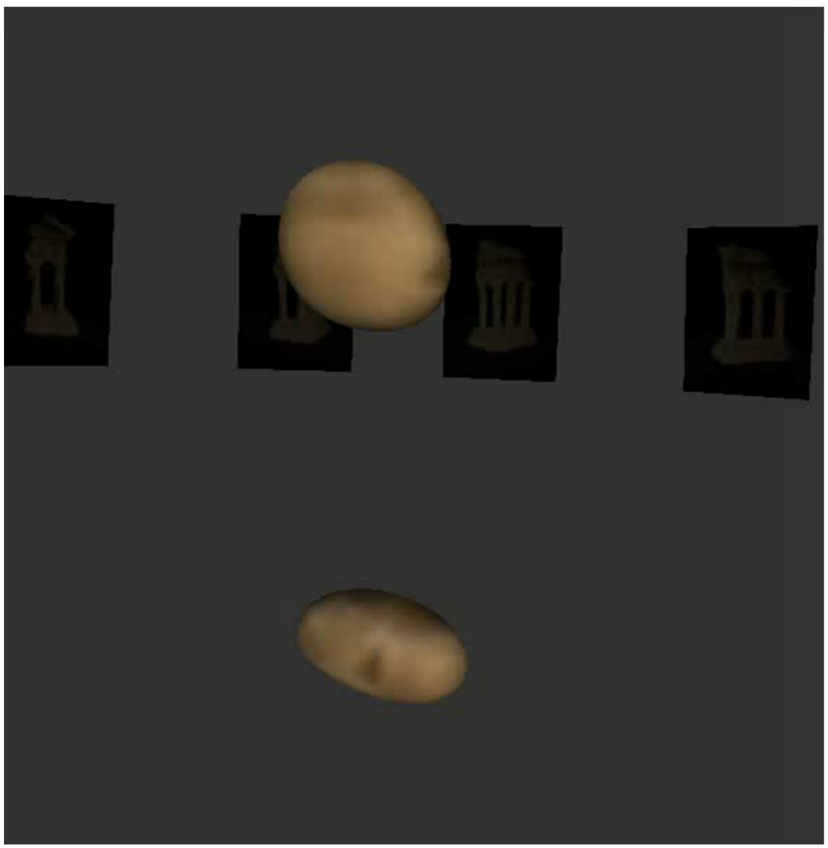

$$v_\Gamma(x) = H(x \bullet n)$$

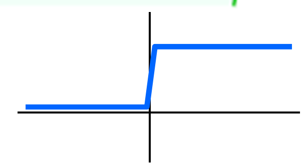

$$H(x \bullet n) = \begin{cases} 1 & x \in (-\pi/2, \pi/2) \\ 0 & x > \pi/2, x < -\pi/2 \\ 1/2 & x \in \{\pi/2, -\pi/2\} \end{cases}$$

Heaviside func.

Dirac func.

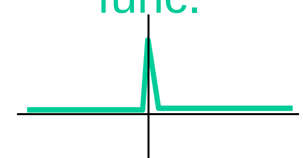

$$\delta(x \bullet n) = \begin{cases} 1 & x \in \{\pi/2, -\pi/2\} \\ 0 & else \end{cases}$$

[Gargallo et al ICCV 07]

Surface discretization

Explicit representation

- **Depth / disparity map**

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

$$E(f) = \int_x \int_y \phi_{data}(x, y, f)\,dxdy + \lambda \int_x \int_y \phi_{reg}(x, y, f, \nabla f)\,dxdy$$

$f_t = \cdots$ (Euler-Lagrange eq.)

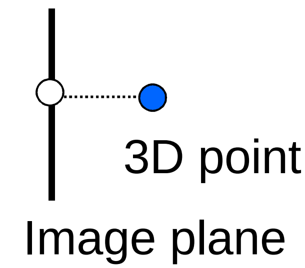

[Robert and Deriche: Dense depth recovery from stereo images, ECAI1992]
[Strecha et al: Dense matching of multiple wide-baseline views ICCV 2003]

- **Depth with respect to plane**

Move points along displacement direction

[Birkbeck et al 2006]

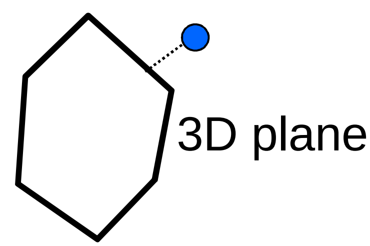

- **Mesh**

$$E(S) \approx \sum_{\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}} \sum_{\{\lambda_1, \lambda_2, \lambda_3\}} \Phi(\lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2 + \lambda_3 \mathbf{v}_3, \lambda_1 \mathbf{n}_1 + \lambda_2 \mathbf{n}_2 + \lambda_3 \mathbf{n}_3)$$

$$\mathbf{v}_t = \partial E / \partial \mathbf{v}$$

[Fua and Leclerc 1993]
[Birkbeck et al ECCV 2006]

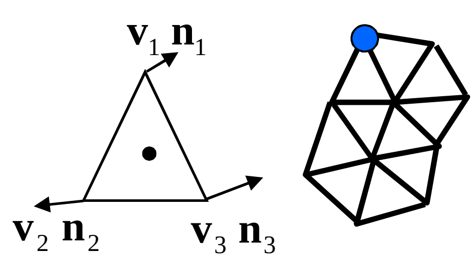

Level set representation

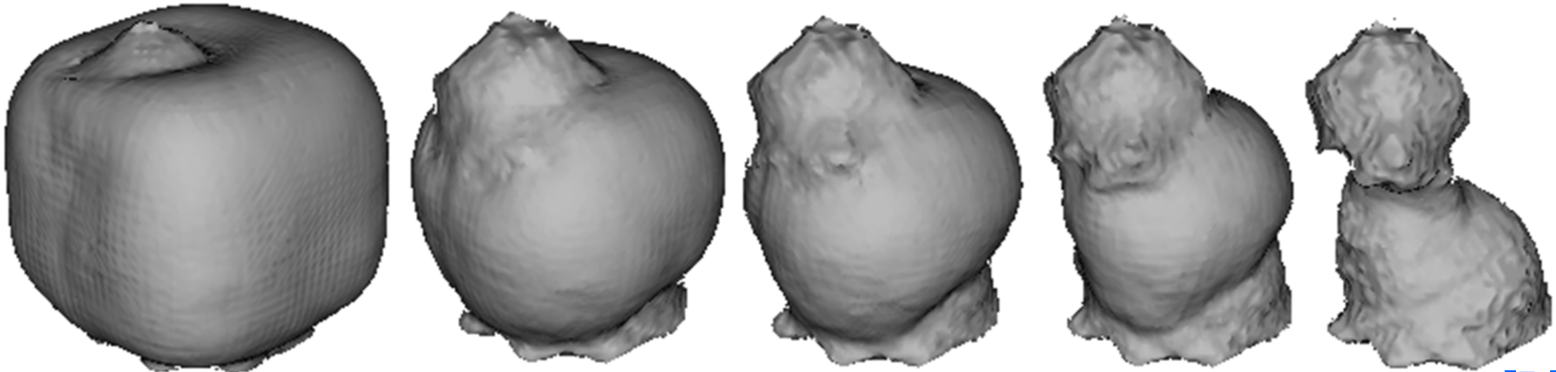

[Birkbeck]

Implicit representation

Surface = Zero level set of a higher dimensional function $\Psi(\mathbf{X},t)$

$\Psi(\mathbf{X}) < 0 \quad \text{for} \quad \mathbf{X} \in \Omega \subset \Re^3$ [Osher and Sethian 88]

$\Psi(\mathbf{X}) > 0 \quad \text{for} \quad \mathbf{X} \notin \overline{\Omega} \text{ (closer of } \Omega)$ [Faugeras Keriven 98] (stereo)

$\Psi(\mathbf{X}) = 0 \quad \text{for} \quad \mathbf{X} \in \partial\Omega = S$ [Soatto Yezzi Jin ICCV 03] (specular refl.)

normal $\mathbf{n} = -\dfrac{\nabla\Psi}{|\nabla\Psi|}$

curvature $H = k = \dfrac{1}{2} \cdot \left(\dfrac{\nabla\Psi}{|\nabla\Psi|} \right)$

Evolution cost function

If $S_t = \Im \mathbf{n} \quad S(0) = S_0$

then $\Psi_t = -\Im \langle \nabla\Psi, \mathbf{n} \rangle \quad \Psi(S(t),t) = 0$

Efficient numerical schemes

[Sethian 96][Osher 02]

Cost functional

Photo-consistency : image cues (SFS,PS,stereo,silhouette)

Where to integrate? : surface
- \+ makes the problem intrinsec
- \+ automatic regularization (multiplicative)
- \- over-smoothing
- \- not account for image discretization

image
- \+ can add a regularizer

Regularization :

linear

quadratic
- \+ improves convergence
- \- penalizes large variation
- \- over-smoothing

non-quadratic, anisotropic $E_{reg} = \int \Phi(|\nabla z|)dxdy$

[Robert, Deriche 96] [Alvarez, Deriche 00] [Strecha 02]

(ex. Nagel-Enkelmann diff. oper.)
- \+ preserves discontinuities

Goto p 66

Surface/depth regularization

Depth map regularization f

energy *Euler Lagrange eq.*

1. <u>Homogenous</u>
Diffusion (heat eq)

$$\left|\nabla f\right|^2 \qquad \Delta f$$

2. <u>Image-based regularization</u> – align depth discontinuities with the image discontinuities

$$g(\left|\nabla I\right|^2)\left|\nabla f\right|^2 \qquad \operatorname{div}\left(g(\left|\nabla I\right|^2)\nabla f\right)$$

$$\nabla f^T D(\nabla I)\nabla f \qquad \operatorname{div}\left(D(\nabla I)\nabla f\right)$$

g – decreasing function

$g(s^2)$->0 when s big (high gradients) – inhibit diffusion

D – tensor $\qquad D = \nabla f \nabla f^T$

3. <u>Depth map regularization</u>
Not smooth across surface discontinuities

$$\phi\left(\left|\nabla f\right|^2\right) \qquad \operatorname{div}\left(\phi'\left(\left|\nabla f\right|^2\right)\nabla f\right) \qquad \phi' = g$$

$$\operatorname{tr}\phi\left(\nabla f \nabla f^T\right) \qquad \operatorname{div}\left(\phi'\left(\underbrace{\nabla f \nabla f^T}_{\mathbf{D}}\right)\nabla f\right)$$

Example regularization

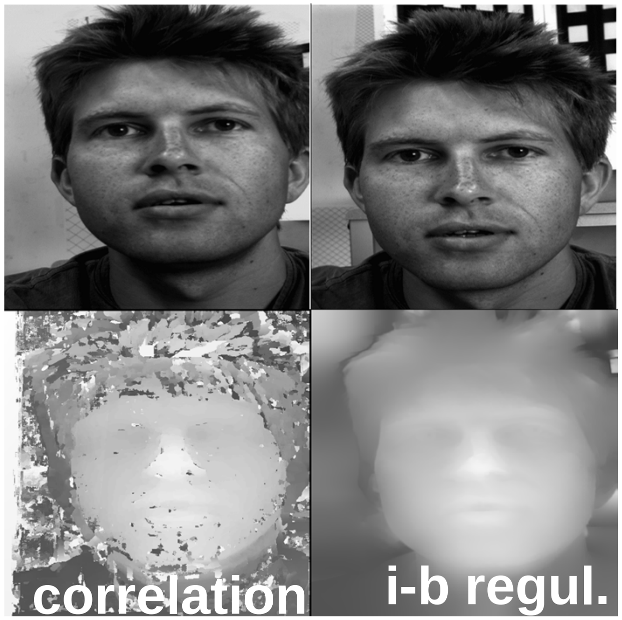

$$g(|\nabla I|^2)|\nabla f|^2 \qquad \mathrm{div}\left(g(|\nabla I|^2)\nabla f\right)$$

$$\nabla f^T D(\nabla I)\nabla f \qquad \mathrm{div}\left(D(\nabla I)\nabla f\right)$$

Image-based
regularization

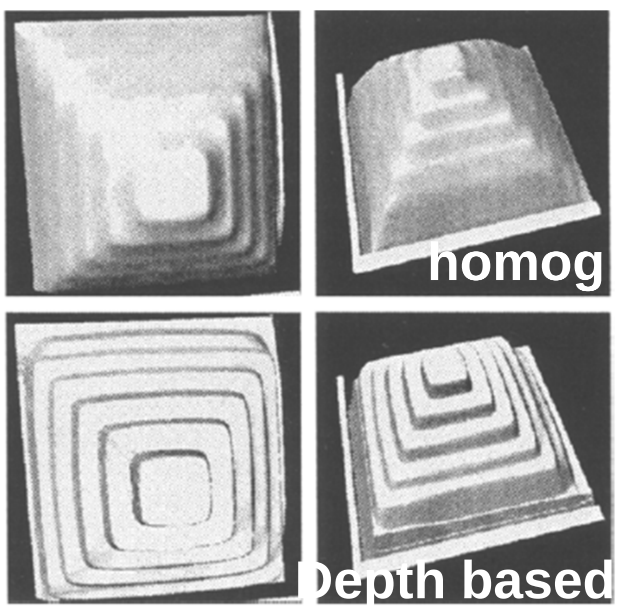

$$\phi\left(|\nabla f|^2\right) \qquad \mathrm{div}\left(\phi'\left(|\nabla f|^2\right)\nabla f\right) \qquad \phi' = g$$

$$\mathrm{tr}\,\phi\left(\nabla f \nabla f^T\right) \qquad \mathrm{div}\left(\phi'\left(\nabla f \nabla f^T\right)\nabla f\right)$$

depth-based
regularization

Example regularization

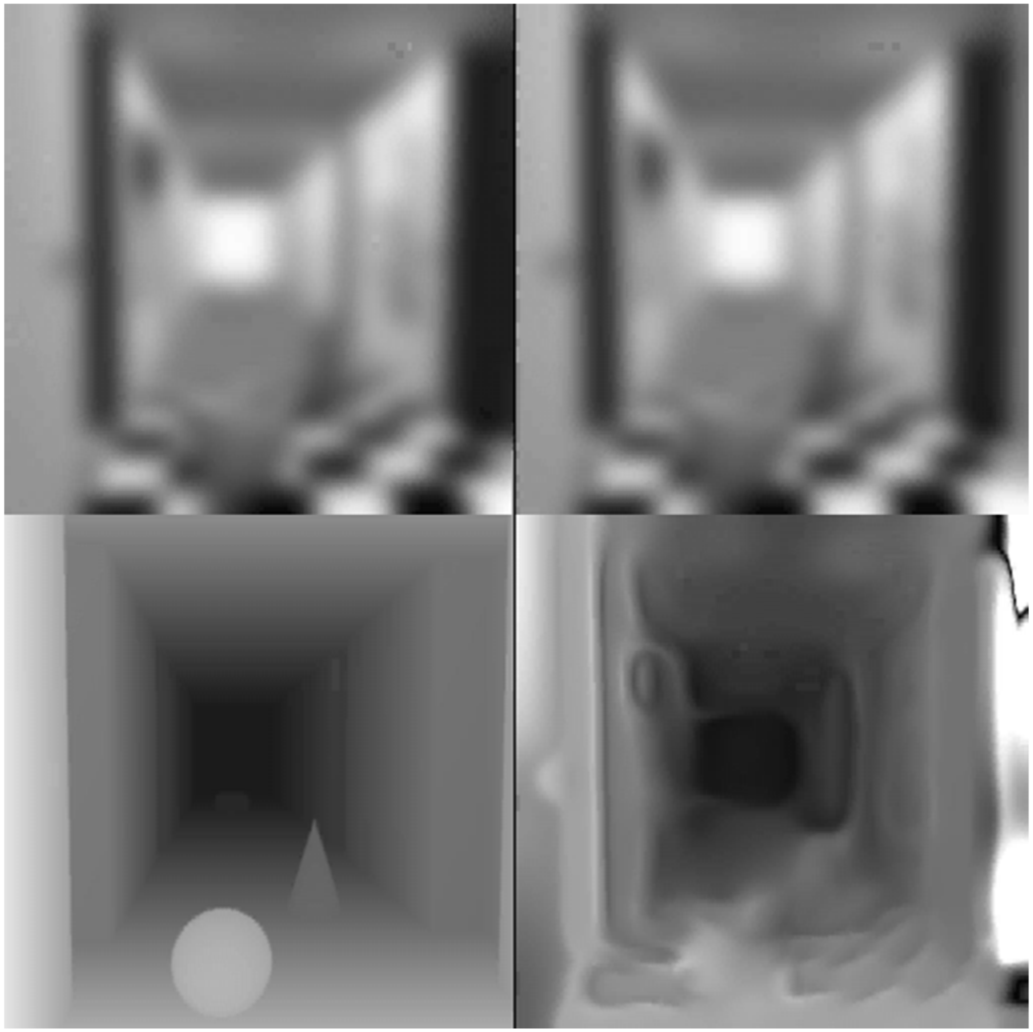

d-b regul　　　i-b regul

Learn more about variational methods

Level Set Methods
S. Osher and R. Fedkiw,
Springer 2003

Jan Erik Solem
PhD thesis,
Malmo Univ.

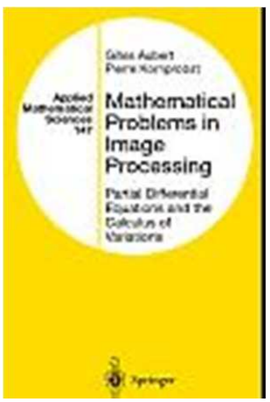

Mathematical Problems in Image Processing
Aubert et al
Springer 2002

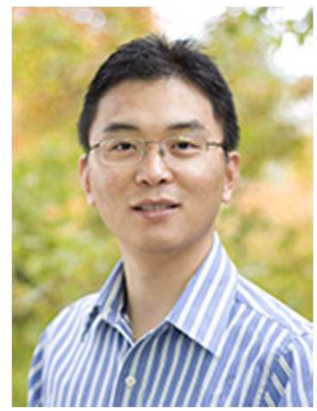

Hailin Jin
PhD thesis,
Washington Univ.

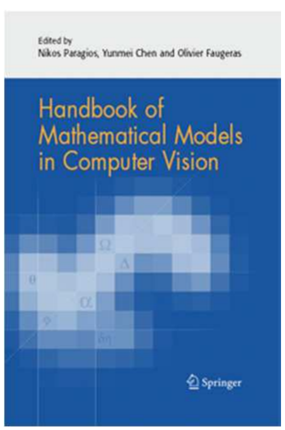

The Handbook of Mathematical Models in Computer Vision
N. Pragios editor
Springer 2005

Graph-cuts and hypersurfaces

- Geometric graph for stereo

[Roy,Cox ICCV 98, IJCV 1999]
[Ishikawa, Geiger ECCV 1998]

- Convex smoothing – global convergence

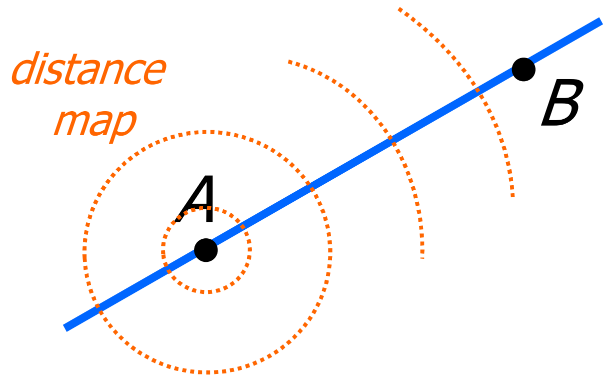

Euclidian metric (ct)

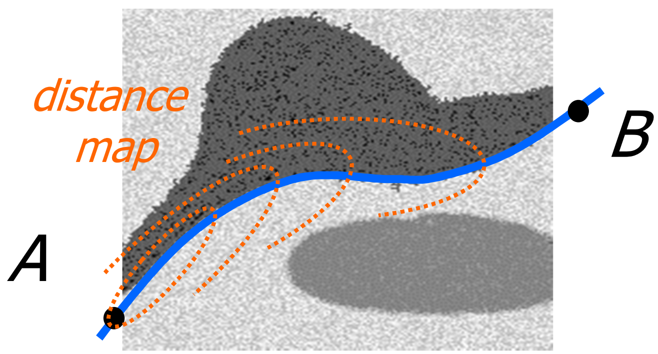

Riemannian metric (varying tensor)

- Connection between cuts and hypersurfaces in continuous spaces

[Boykov, Kolmogorov: Computing geodesics and minimal surfaces via graph cuts, ICCV 2003]

Show how to build a grid graph and sets the weights such that the cost of cuts is arbitrarily close to the area of corresponding surface for any anisotropic Riemannian metric.

Graph cut to find globally minimum surfaces (like level sets) under arbitrary Riemannian metric.

Minimal surfaces and graph cut

[Boykov, Kolmogorov ICCV03,05]

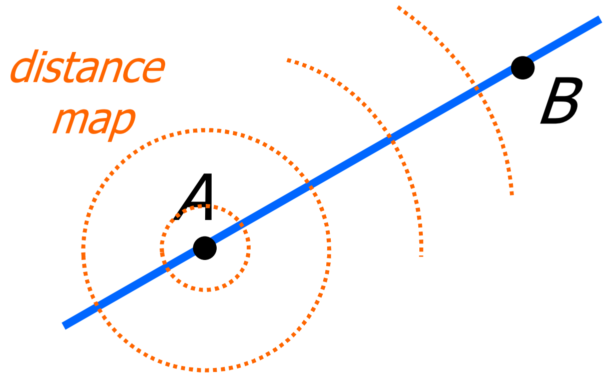

Euclidian metric (ct)

$$E = \int_S dA$$

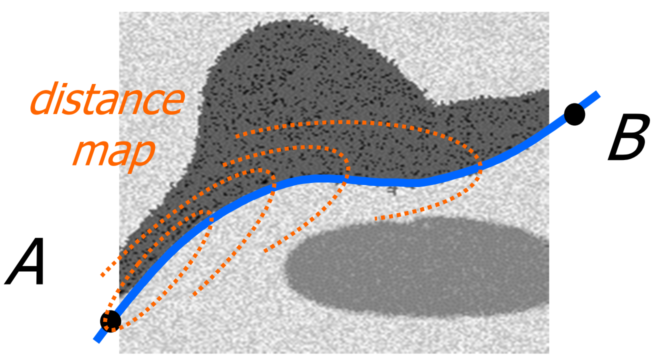

Riemannian metric (varying tensor)

$$E = \int_S \Phi(\mathbf{X})\, dA$$

multi view stereo formulation

Can the minimal surface energy be minimized with graph cut ?

Cost of a cut

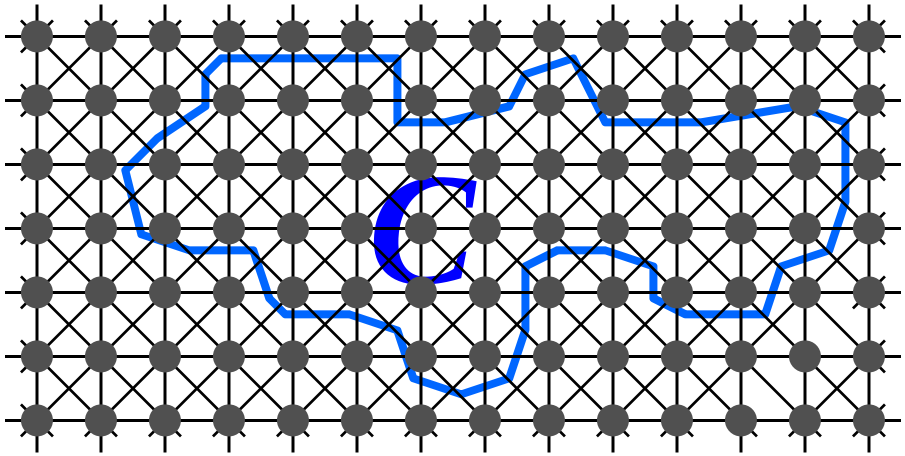

$$\| C \| = \sum_{e \in C} | e |$$

- Cost of a cut can be interpreted as a geometric "length" (in 2D) or "area" (in 3D) of the corresponding contour/surface.

- *Cut metric* is determined by the graph *topology* and by *edge weights*.

Riemanian metric

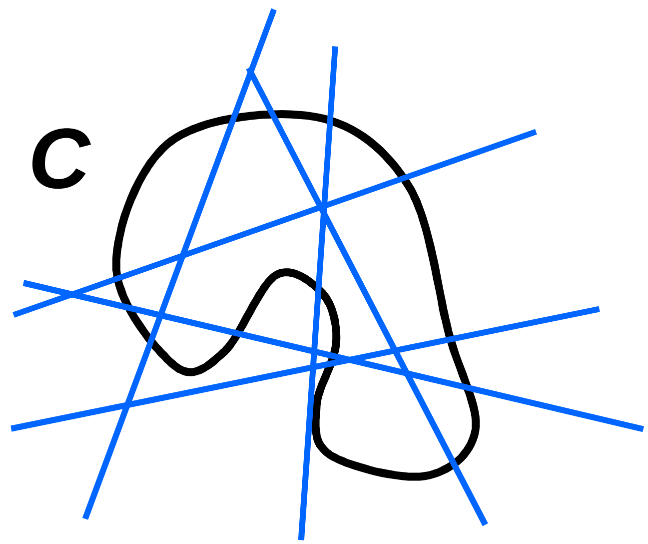

Euclidean length of ***C***
Cauchy-Crofton formula

$$\boxed{\| C \|_{\varepsilon} \;=\; \frac{1}{2} \int n_L \cdot d\rho \cdot d\phi}$$

the number of times
line *L* intersects *C*

Cut Metric on grids can approximate Euclidean Metric

Graph nodes are imbedded in R2 in a grid-like fashion

Edges of any regular neighborhood system generate families of lines

{—, /, |, \}

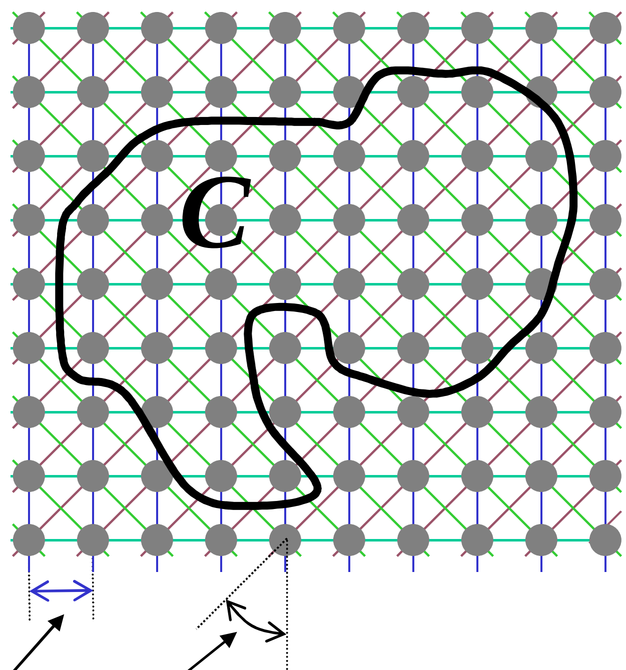

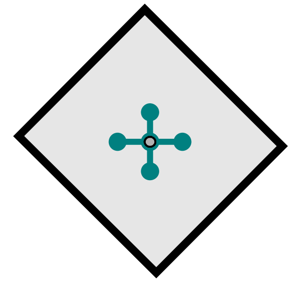

4-neighborhoods (*Manhattan* metric)

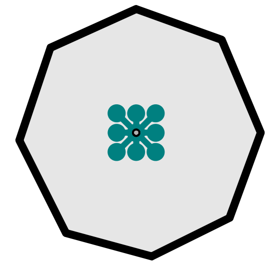

8-neighborhoods

$$\| C \|_{\varepsilon} \approx \tfrac{1}{2} \sum_{k} n_k \cdot \Delta\rho_k \cdot \Delta\phi_k = \| C \|_{gc}$$

Euclidean length

the number of edges of family k intersecting *C*

graph cut cost for edge weights:

$$\boxed{w_k = \frac{\Delta\rho_k \cdot \Delta\phi_k}{2}}$$

Example [Boykov]

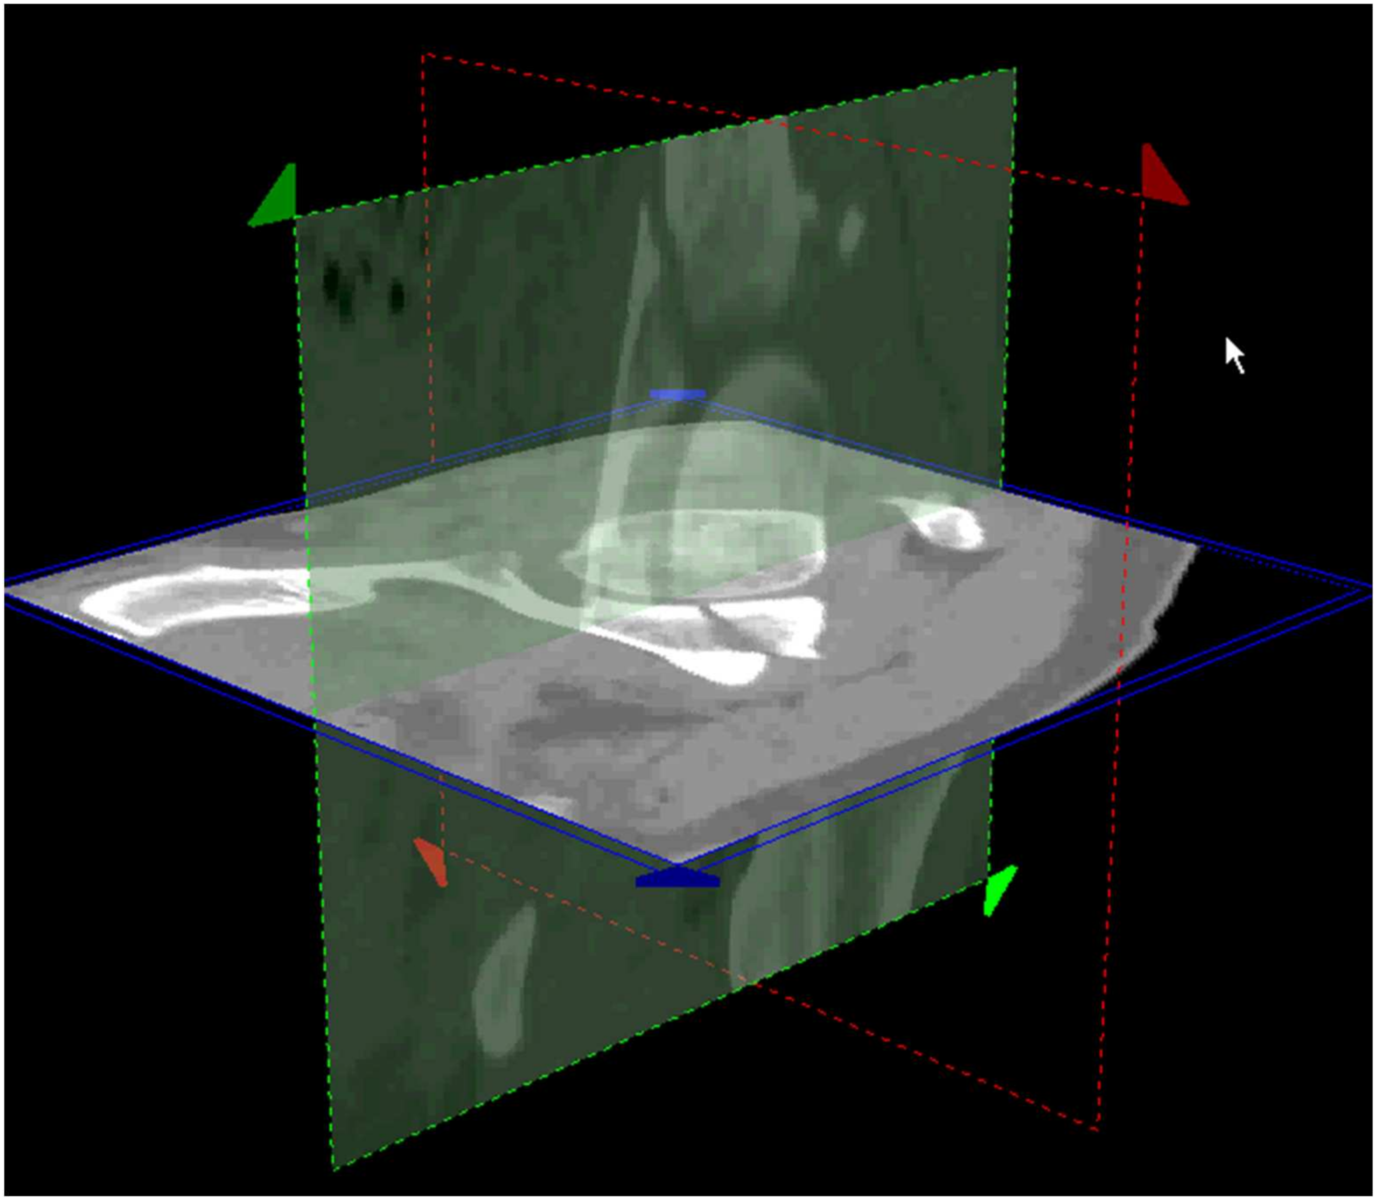

Summary: representations

Image centered

- Depth/disparity map

Object centered

Implicit (level sets) | Mesh | Voxels

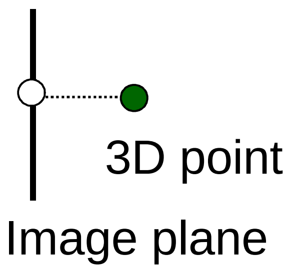

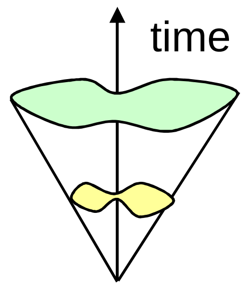

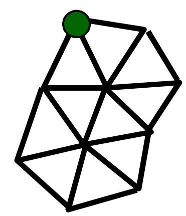

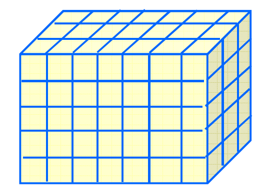

	Depth/disparity map	Implicit (level sets)	Mesh	Voxels
+	▪Natural extension of SFS, PS, stereo ▪Strong min with graph cuts	▪Handles topological changes ▪Implied normals ▪Visibility	▪Graphics based ▪Implied normals ▪Visibility	▪Handles large structures ▪Arbitrary topology
-	▪Limited resolution ▪Partial obj. reconstr ▪Viewpoint dependent	▪Too smooth ▪Slow ▪Closed surfaces	▪Topological changes	▪Ignores regularization (sensitive to noise) ▪Limited resolution ▪Occlusion handling

Discrete vs. continuous

Level Sets	Graph Cuts
variational optimization method for fairly general continuous energies	combinatorial optimization for a restricted class of energies [e.g. KZ'02]
finds a local minimum near given initial solution	finds a global minimum for a given set of boundary conditions
numerical stability has to be carefully addressed [Osher&Sethian'88] *continuous formulation -> "finite differences"*	numerical stability is not an issue *discrete formulation ->min-cut algorithms*
anisotropic metrics are harder to deal with (e.g. slower)	anisotropic Riemannian metrics are as easy as isotropic ones

Gradient descent method VS. Global minimization tool (restricted class of energies)

[Boykov CVPRTut 2005]

A complete system

Neil Birkbeck

Motivation

method that works for objects with general reflectance

- textured and uniform regions ⇨ shading + texture
  light variation (multiview PS)
- specular materials ⇨ general BRDF – parametric

System

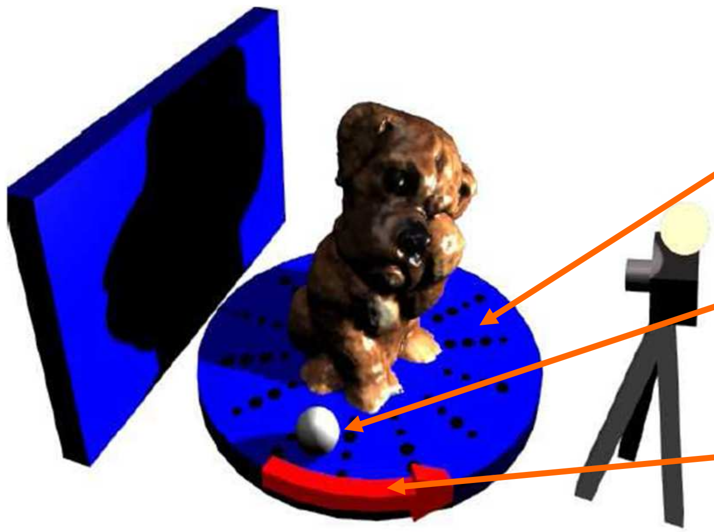

Camera calibration:
pattern

Light calibration:
specular white sphere

Light variation:
rotating table, fixed cam.

Cost functional

Cost functional

$$E = \int_S \Phi(\mathbf{X},\mathbf{n})\,dA = \int_u \int_v \Phi(\mathbf{X},\mathbf{n}) \left\| \mathbf{X}_u \times \mathbf{X}_v \right\| du\,dv$$

$$\mathbf{n} = \frac{\mathbf{X}_u \times \mathbf{X}_v}{\left\| \mathbf{X}_u \times \mathbf{X}_v \right\|}$$

Per-point cost function

$$\Phi(\mathbf{X},\mathbf{n}) = \sum_i h(\mathbf{X},P_i) \left\| I_i(P_i(\mathbf{X})) - R(\mathbf{X},\mathbf{n},\mathbf{L}_i) \right\|$$

Visibility+sampling

reflectance

camera proj.

image

light

1. Lambertial reflectance

$$R_{lamb} = \left(\langle \mathbf{n}, \mathbf{l}_i \rangle l_i + a_i \right) k_{d,\mathbf{X}}$$

light dir.

light color

ambient color

BRDF albedo

2. Non-Lambertian reflectance

BRDF = Phong

<u>In practice</u> : filter specular highlights

fit full reflectance only at the end

Surface discretization

Surface = triangles,
move vertices

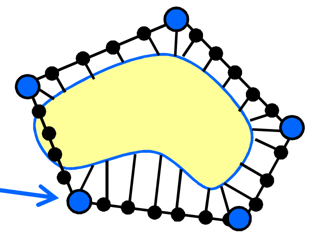

$$\mathbf{X} = \lambda_1 \boxed{\mathbf{v}_1} + \lambda_2 \boxed{\mathbf{v}_2} + \lambda_3 \boxed{\mathbf{v}_3}$$

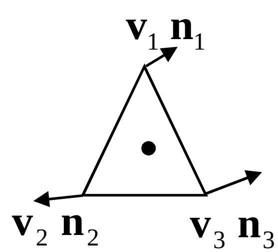

Cost function

$$E(S) \approx \sum_{\{\mathbf{v}_1,\mathbf{v}_2,\mathbf{v}_3\}} \sum_{\{\lambda_1,\lambda_2,\lambda_3\}} \Phi(\lambda_1\mathbf{v}_1 + \lambda_2\mathbf{v}_2 + \lambda_3\mathbf{v}_3, \lambda_1\mathbf{n}_1 + \lambda_2\mathbf{n}_2 + \lambda_3\mathbf{n}_3)$$

Normals interpolated

$$\mathbf{n} = \lambda_1\mathbf{n}_1 + \lambda_2\mathbf{n}_2 + \lambda_3\mathbf{n}_3$$

$$\mathbf{n}_j = \sum_{j,k,i\in\Delta(\mathbf{v}_j)} (\mathbf{v}_k - \mathbf{v}_j)\times(\mathbf{v}_i - \mathbf{v}_j)$$

Albedo implied by the shape (closed form solution)

(knowing light dir+color)

Visibility/shadows Z buffering

Surface discretization

Regularizer mean curvature

$$\frac{\partial S}{\partial t} = \left(2\Phi H - \langle \nabla\Phi, \mathbf{n} \rangle\right)\mathbf{n}$$

Gradient finite differences

Initial shape = visual hull

Mesh handles topological changes

Multi-resolution (image pyramid)

Results : refinement

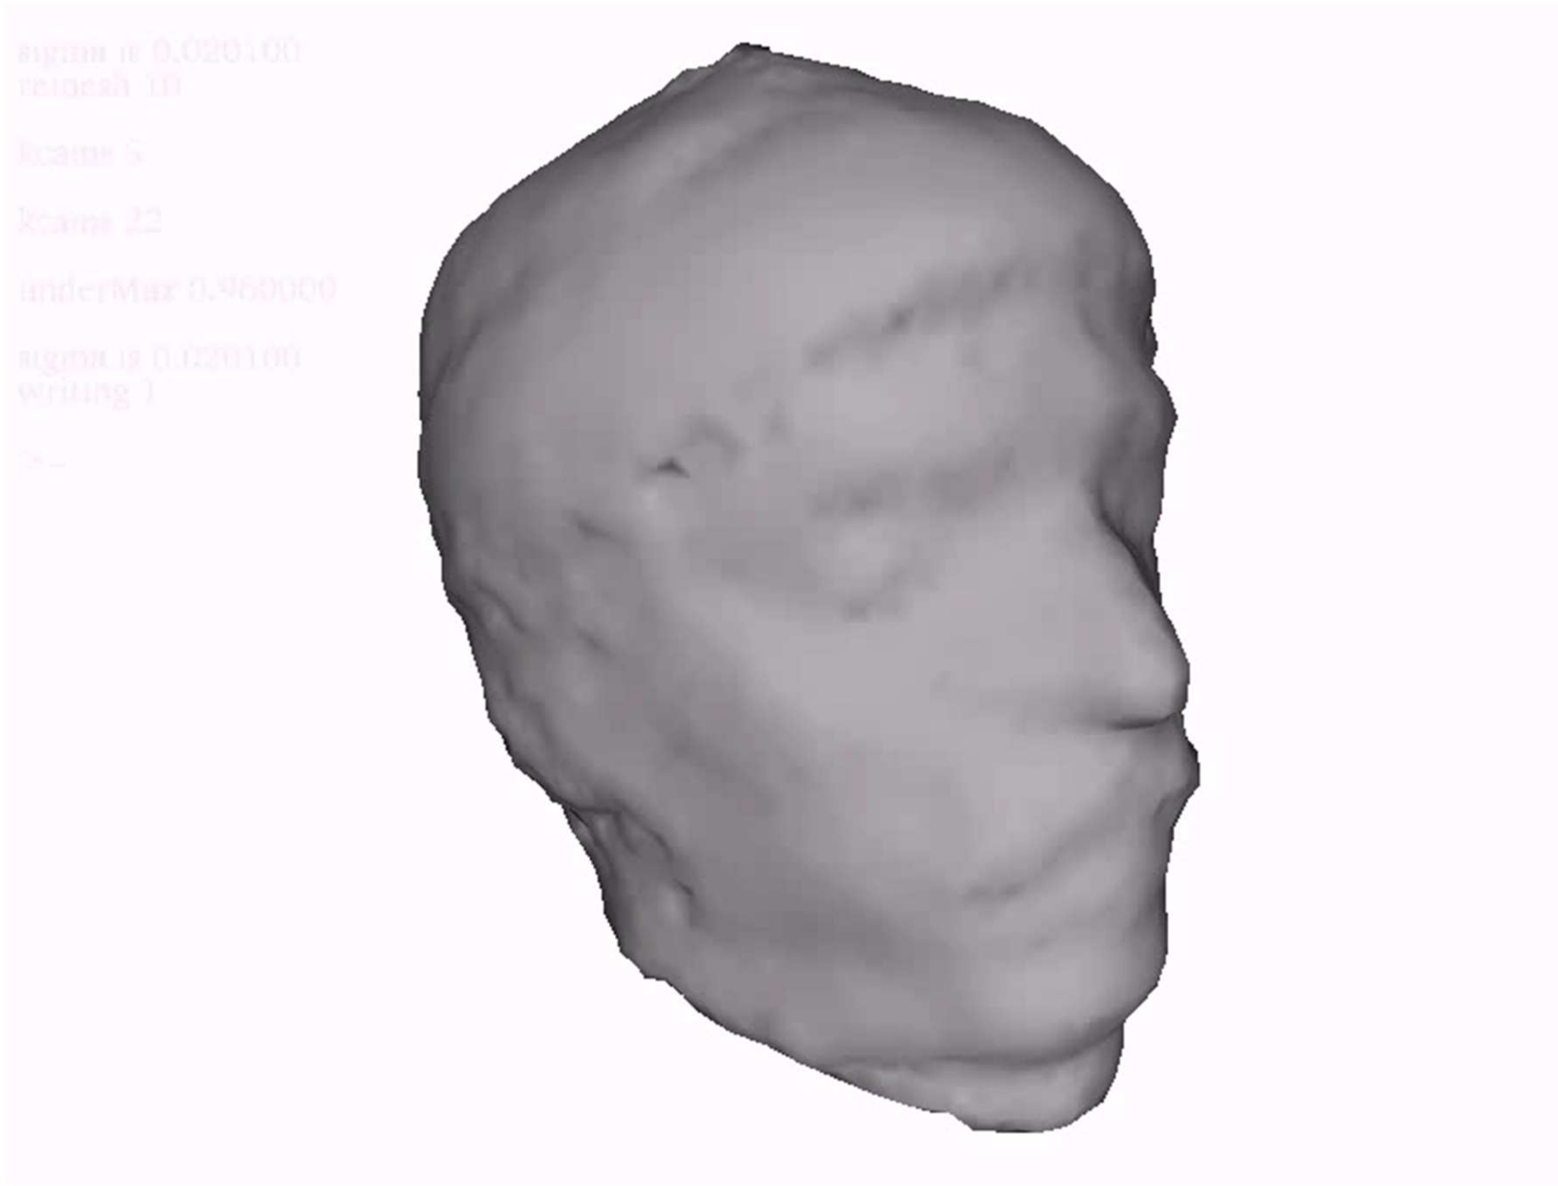

Results : shape + reflectance

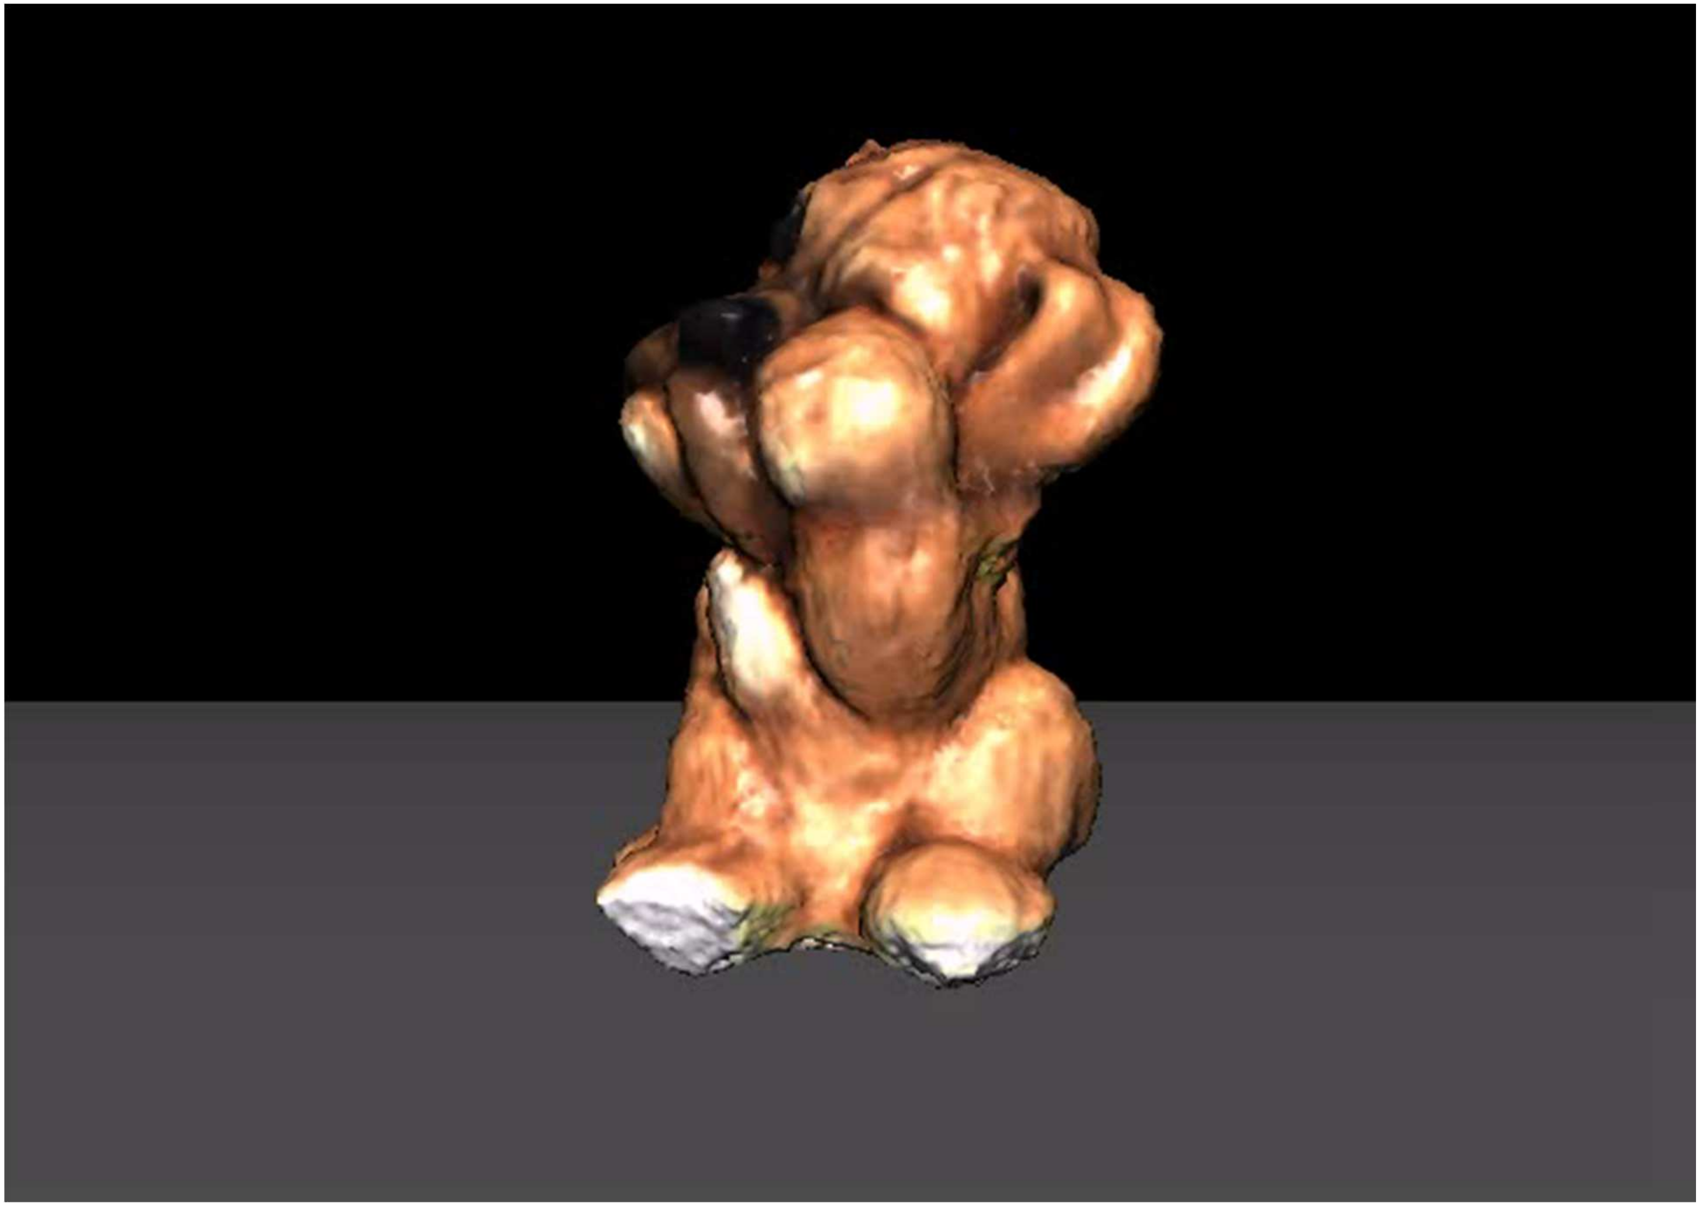

Improvements

- Better light model – no need for calibration
- Model background ? no need to extract silhouette (foreground) – Mumford-Shah functional
- Discontinuities – anisotropic regularization; discontinuous mesh
- Incorporate noise ?

Capture system

- Shape from silhouette
- Dynamic texture