ECCV 2006 tutorial on

Graph Cuts vs. Level Sets

part I

Basics of Graph Cuts

Yuri Boykov
University of Western Ontario

Daniel Cremers
University of Bonn

Vladimir Kolmogorov
University College London

Graph Cuts versus Level Sets

- Part I: Basics of *graph cuts*
- Part II: Basics of *level-sets*
- Part III: Connecting *graph cuts* and *level-sets*
- Part IV: Global vs. local optimization algorithms

Graph Cuts versus Level Sets

■ Part I: Basics of *graph cuts*

- Main idea for min-cut/max-flow methods and applications
- Implicit and explicit representation of boundaries
- Graph cut energy (different views)
  - submodularity, geometric functionals, posterior energy (MRF)
- Extensions to multi-label problems
  - convex and non-convex (robust) interactions, a-expansions

1D Graph cut ⇔ shortest path on a graph

Example:
find the shortest closed contour in a given domain of a graph

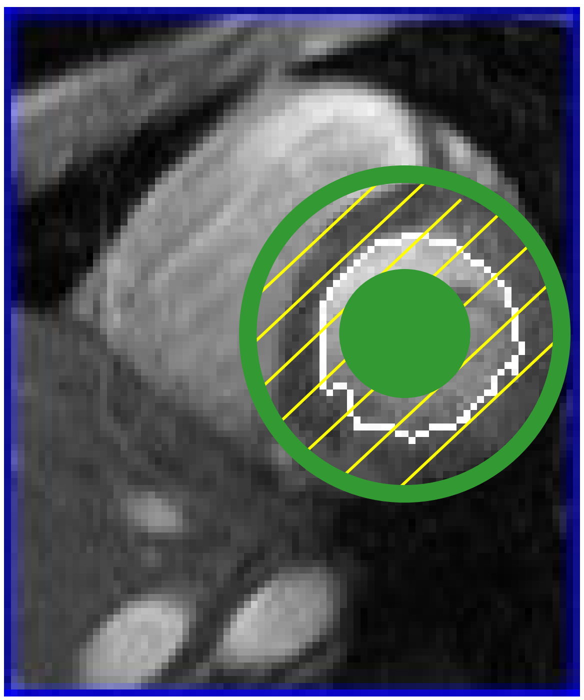

Shortest paths approach

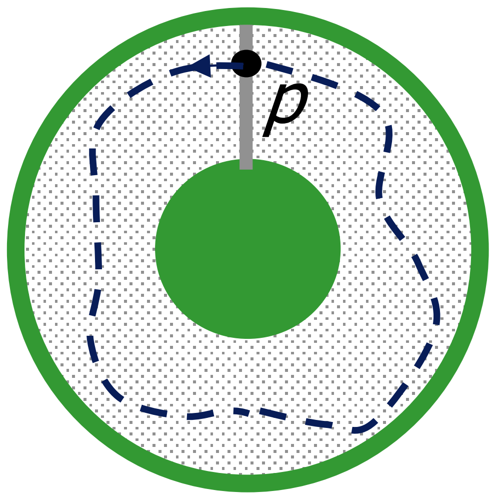

Compute the *shortest path* *p ->p* for a point *p*. Repeat for all points on the gray line. Then choose the optimal contour.

Graph Cuts approach

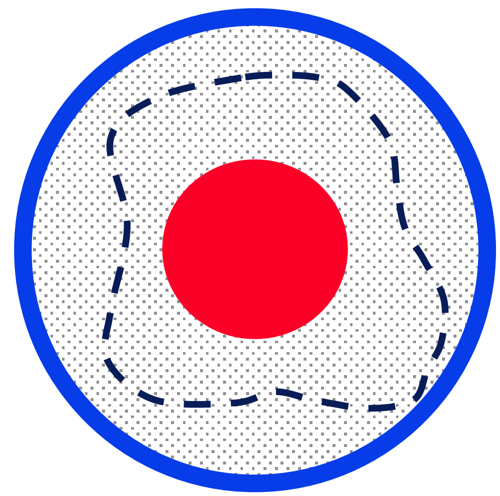

Compute the *minimum cut* that separates red region from blue region

Graph cuts for optimal boundary detection
(simple example à la Boykov&Jolly, ICCV'01)

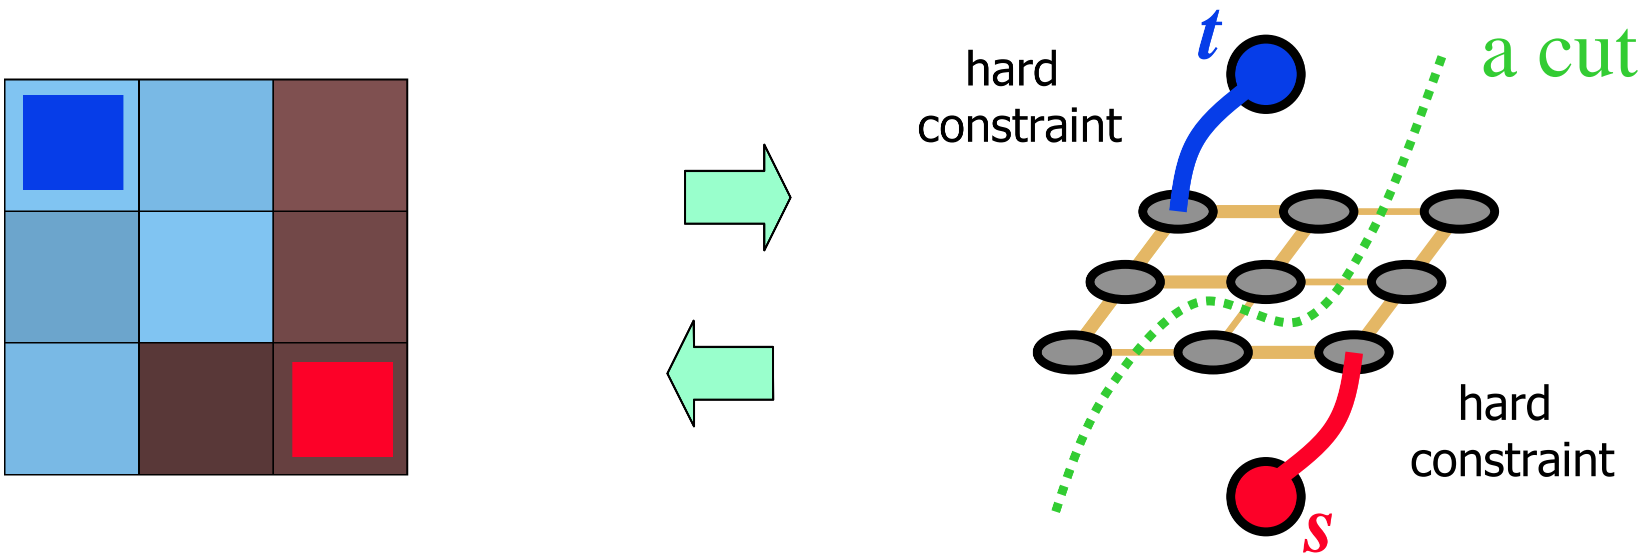

$$w_{pq} = \exp\left\{-\frac{\Delta I_{pq}}{2\sigma^2}\right\}$$

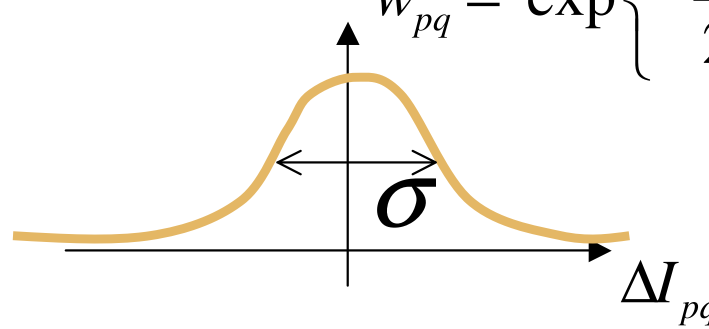

Minimum cost cut can be computed in polynomial time

(max-flow/min-cut algorithms)

Minimum *s-t* cuts algorithms

- Augmenting paths [Ford & Fulkerson, 1962]
- Push-relabel [Goldberg-Tarjan, 1986]

"Augmenting Paths"

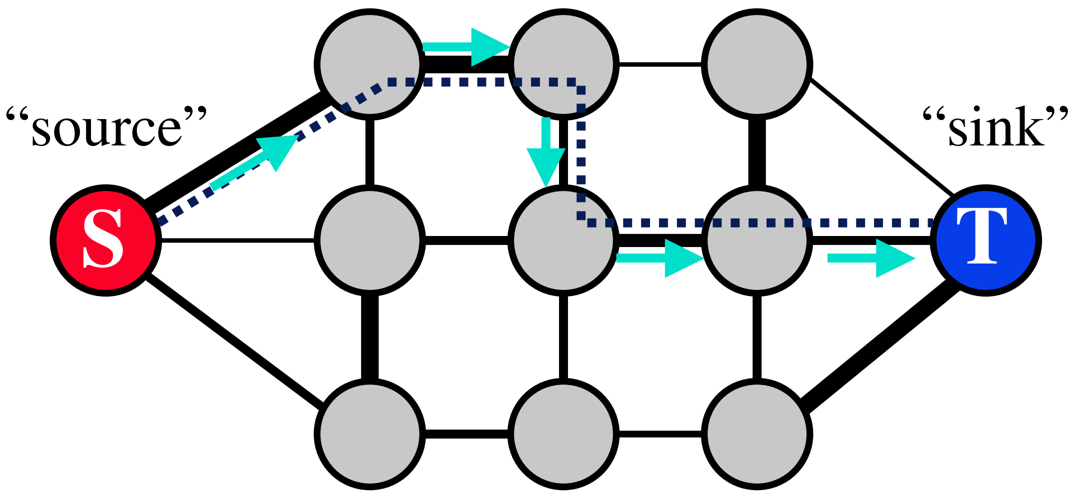

A graph with two terminals

- Find a path from S to T along non-saturated edges
- Increase flow along this path until some edge saturates

"Augmenting Paths"

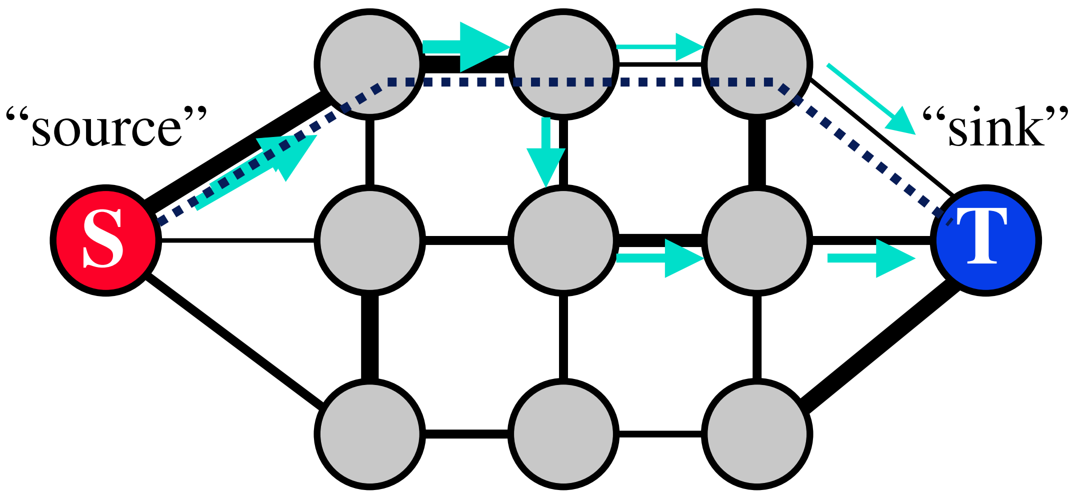

A graph with two terminals

- Find a path from S to T along non-saturated edges
- Increase flow along this path until some edge saturates
- Find next path…
- Increase flow…

"Augmenting Paths"

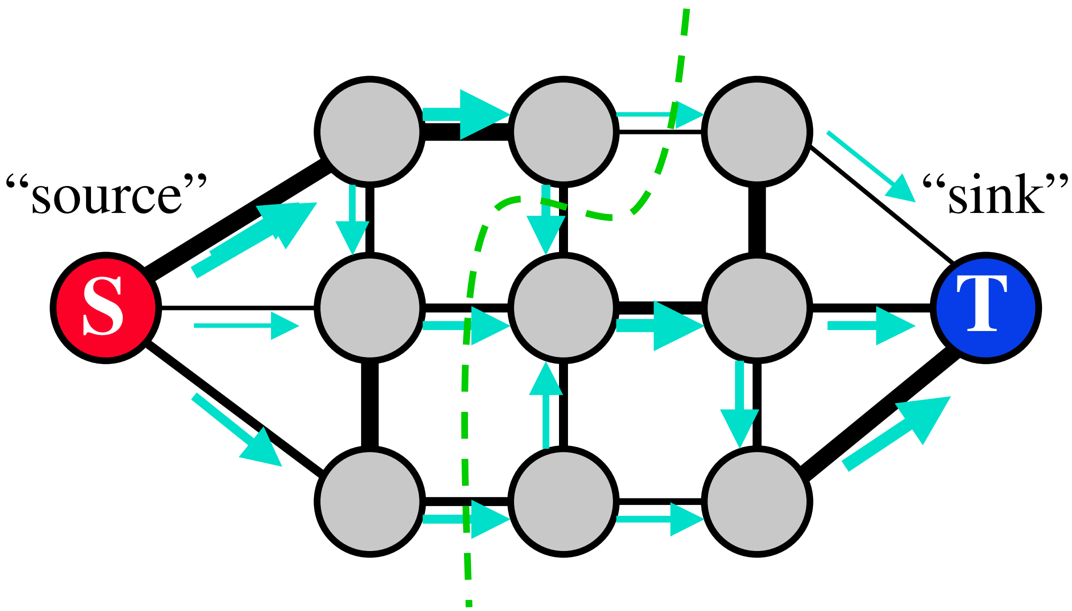

A graph with two terminals

MAX FLOW ⇔ MIN CUT

- Find a path from S to T along non-saturated edges
- Increase flow along this path until some edge saturates

Iterate until …
all paths from S to T have at least one saturated edge

Optimal boundary in 2D

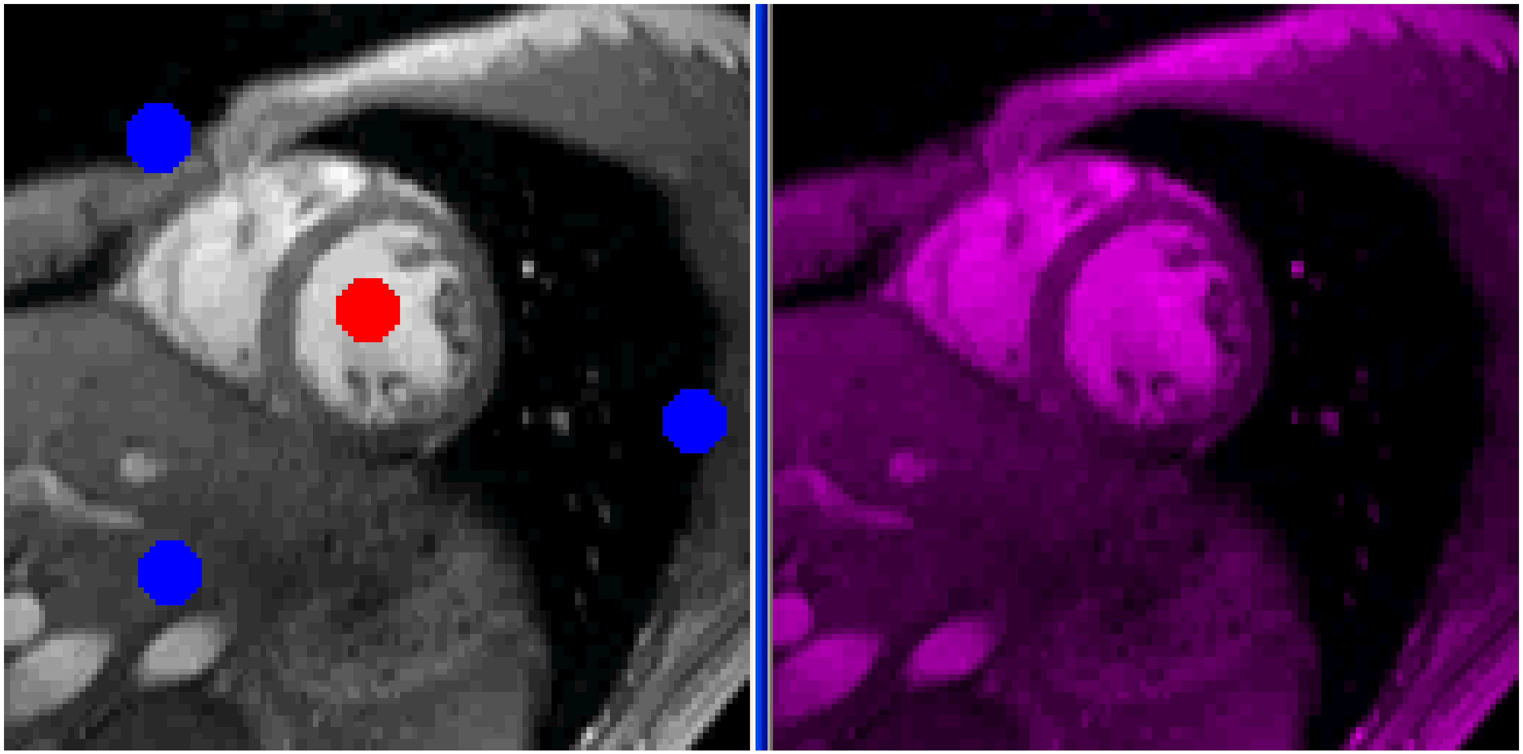

“max-flow = min-cut”

Optimal boundary in 3D

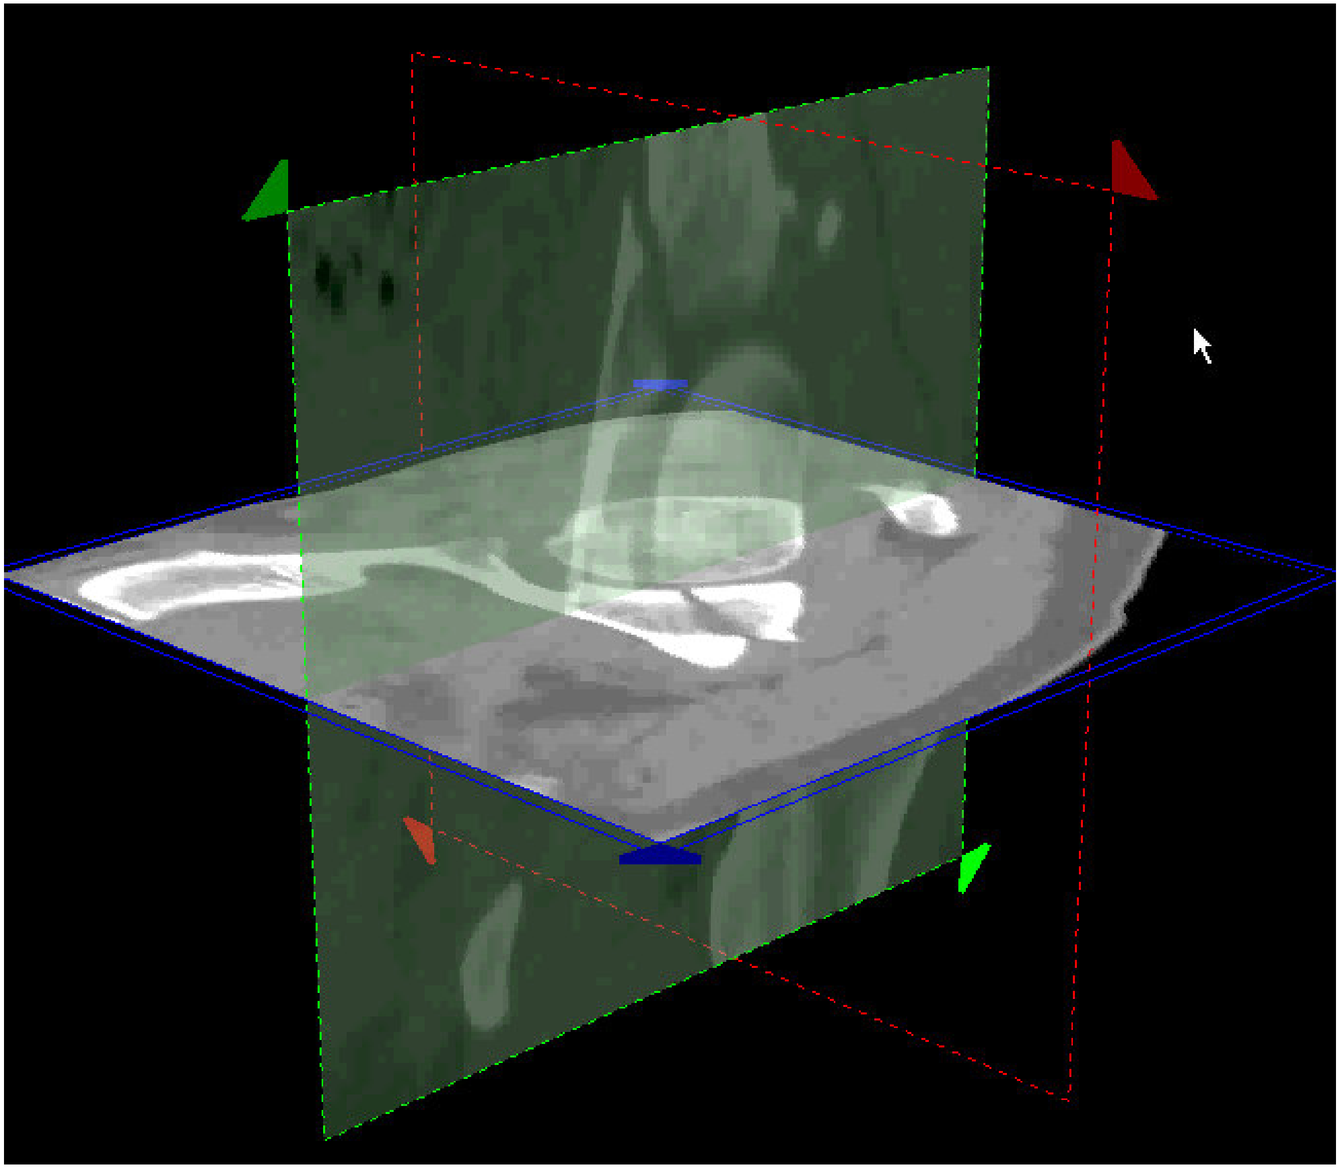

3D bone segmentation (real time screen capture)

Graph cuts applied to multi-view reconstruction

Calibrated images of Lambertian scene

3D model of scene

CVPR'05 slides from Vogiatzis, Torr, Cippola

Graph cuts applied to multi-view reconstruction

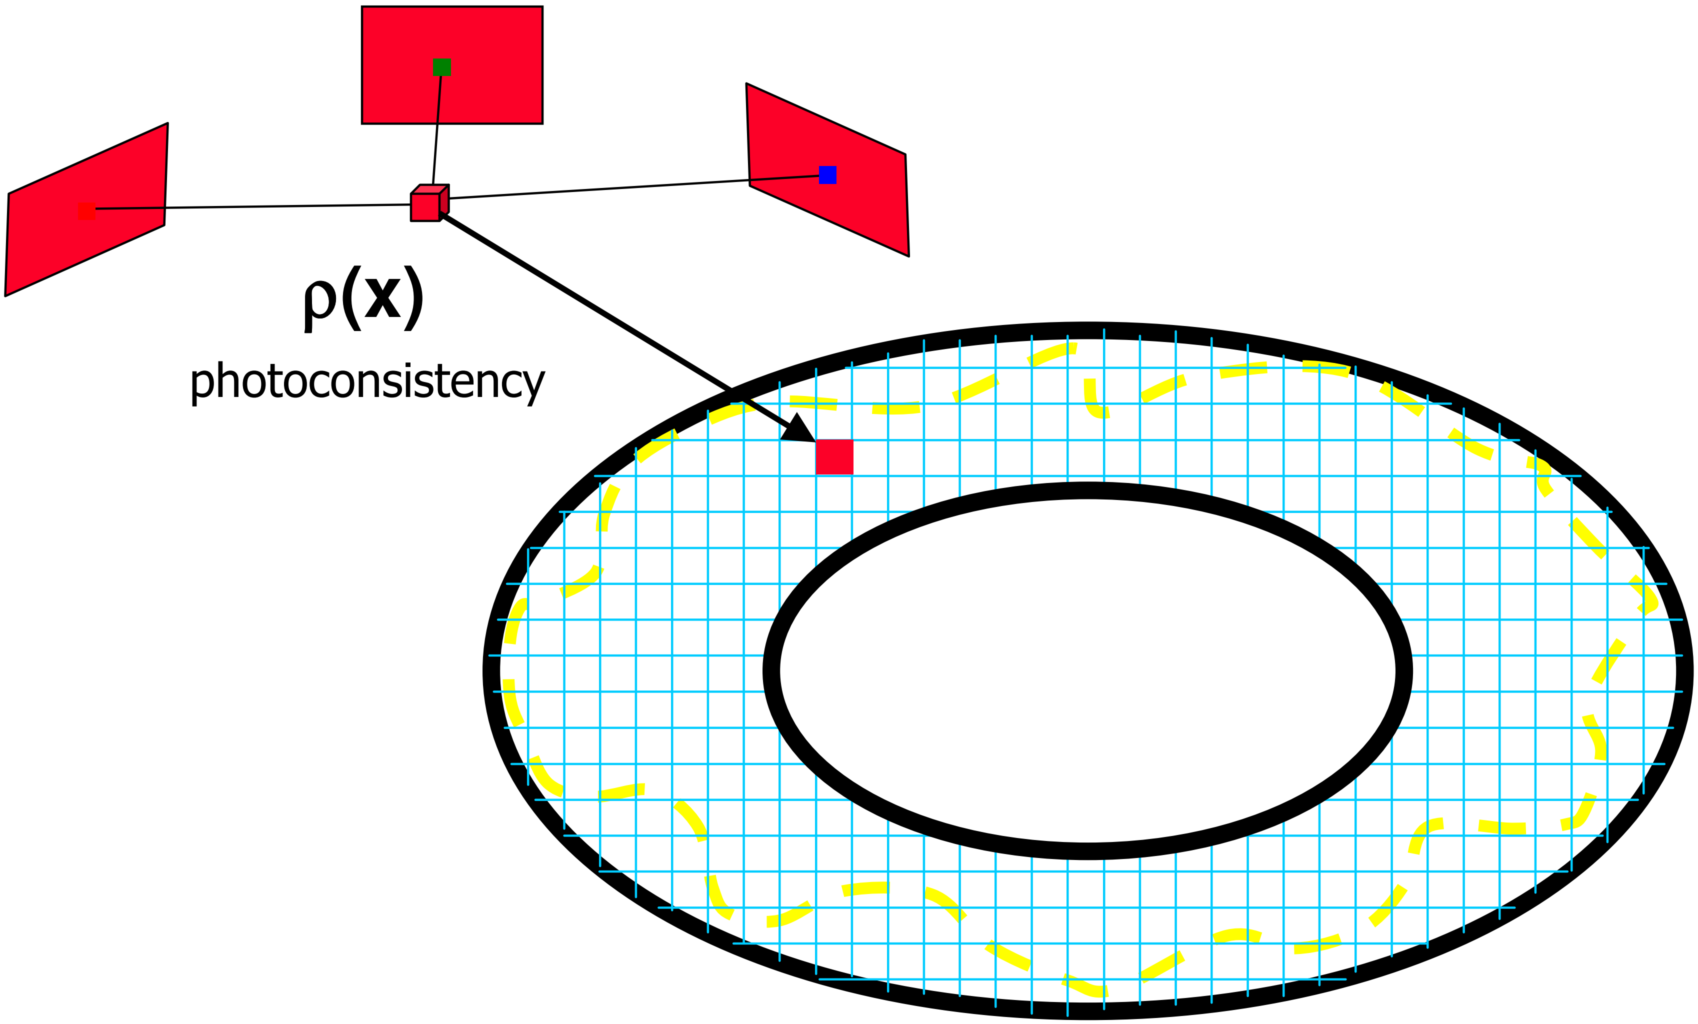

CVPR'05 slides from Vogiatzis, Torr, Cippola

Estimating photoconsistency in a narrow bad

- Occlusion

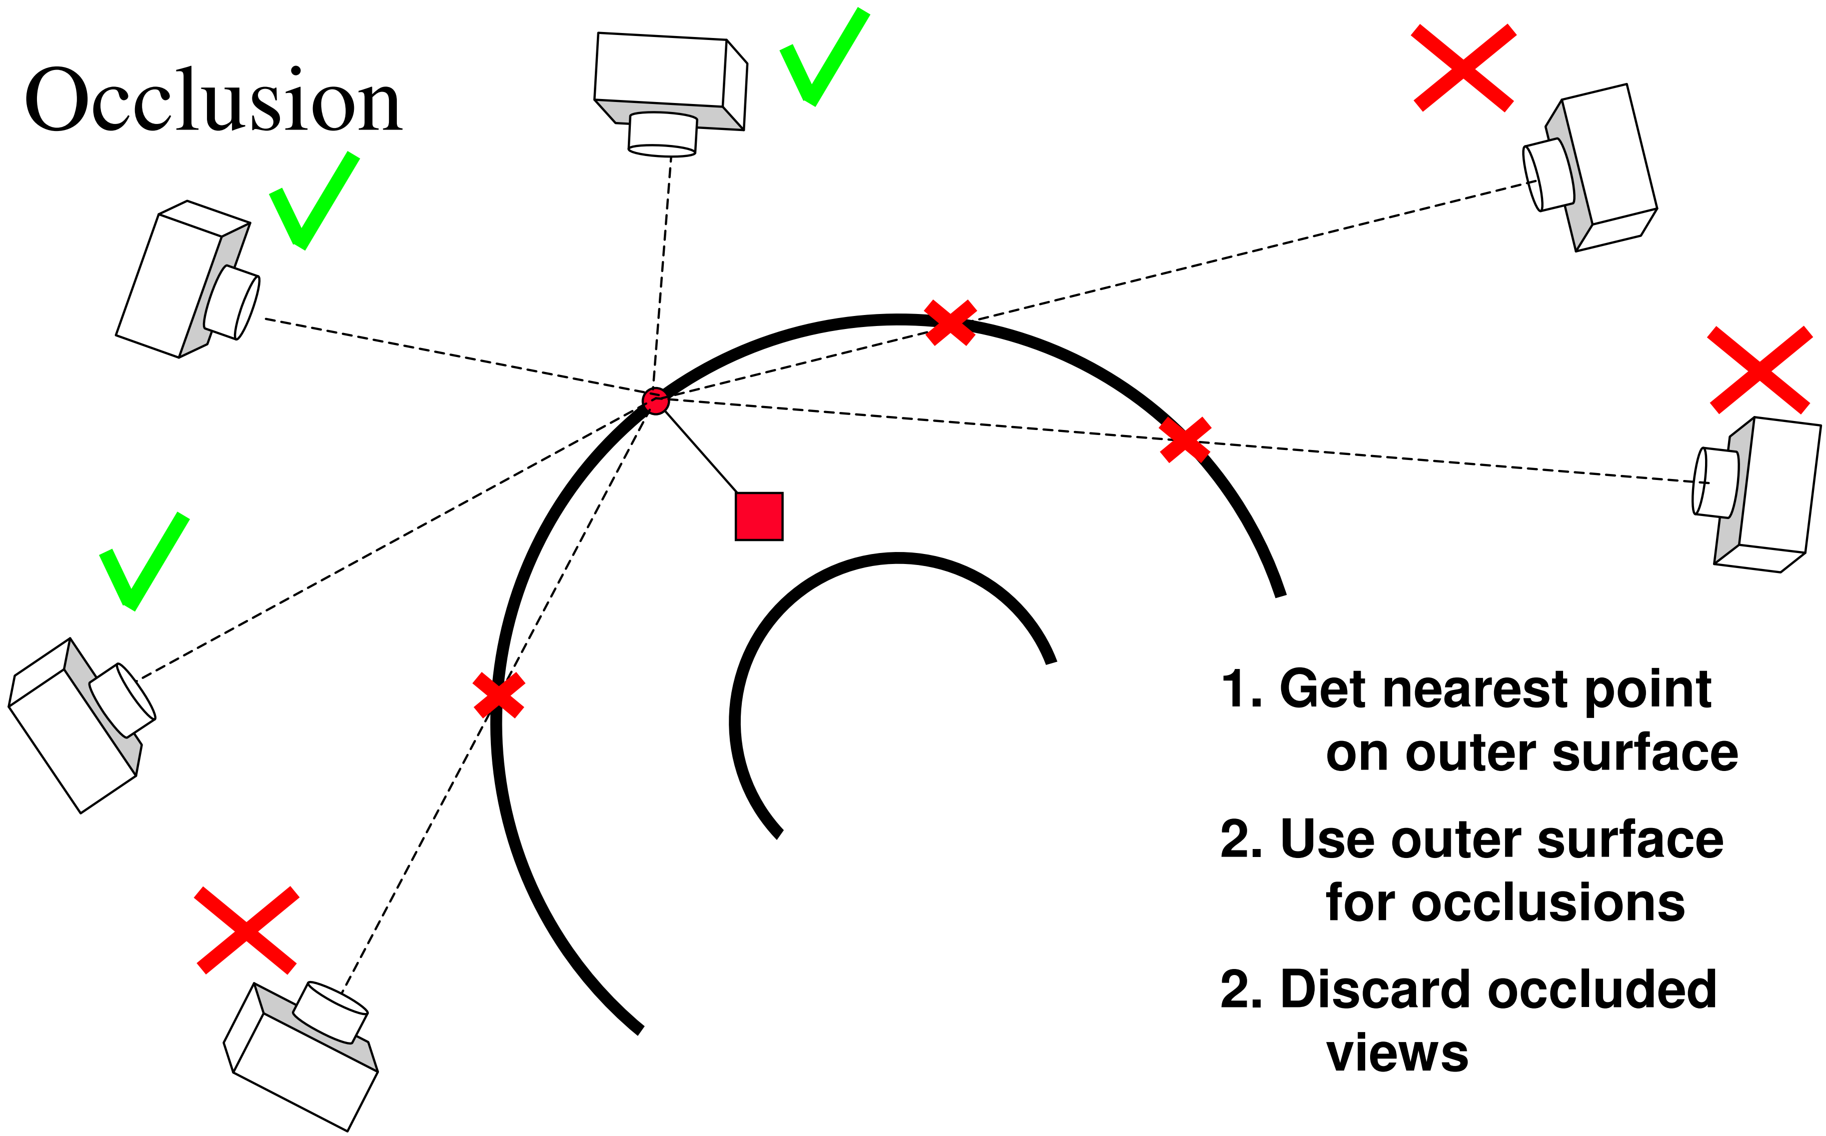

1. Get nearest point on outer surface
2. Use outer surface for occlusions
2. Discard occluded views

CVPR'05 slides from Vogiatzis, Torr, Cippola

Graph cuts applied to multi-view reconstruction

The cost of the cut integrates photoconsistency over the whole space

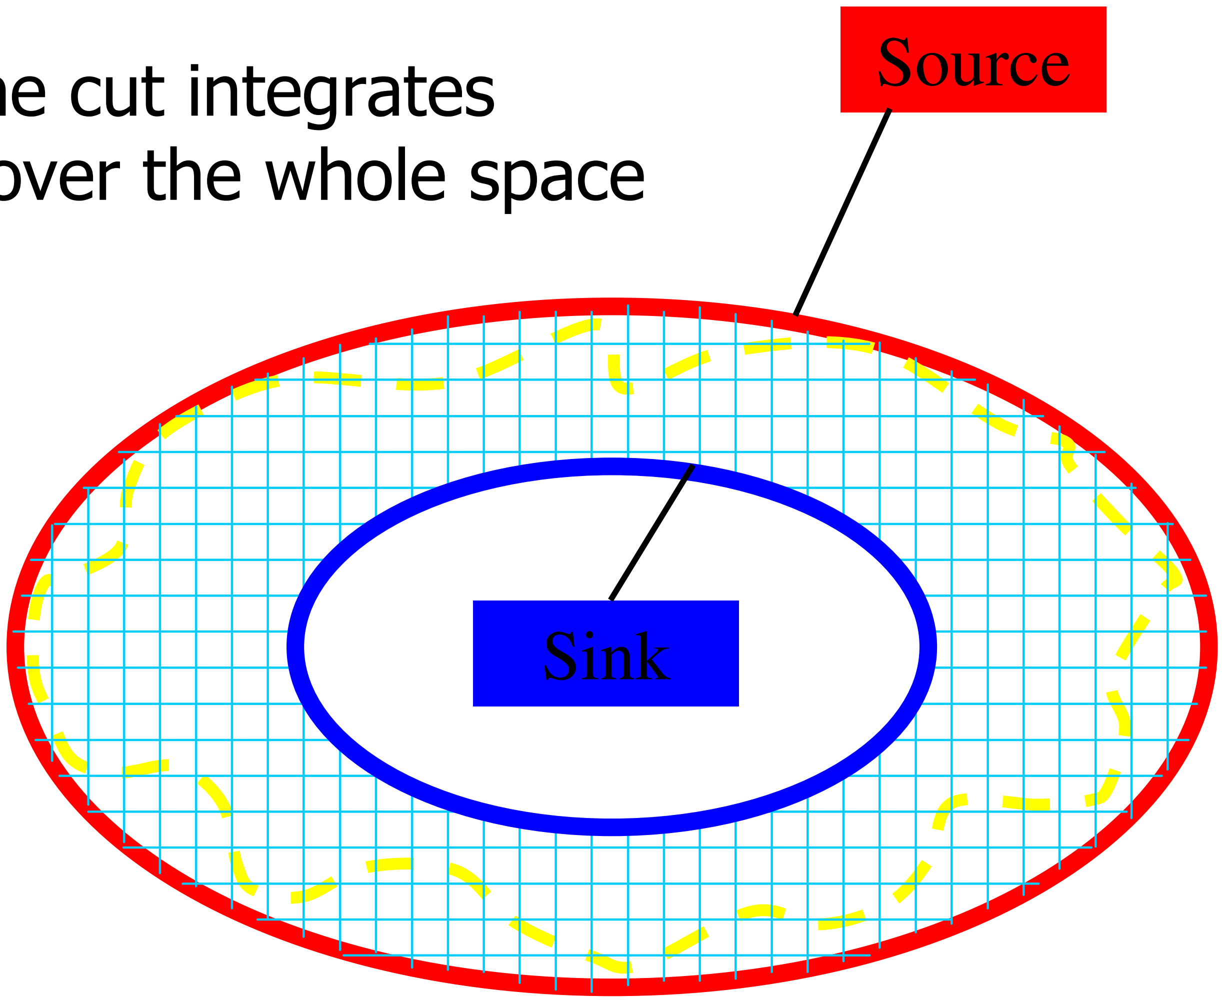

CVPR'05 slides from Vogiatzis, Torr, Cippola

Graph cuts applied to multi-view reconstruction

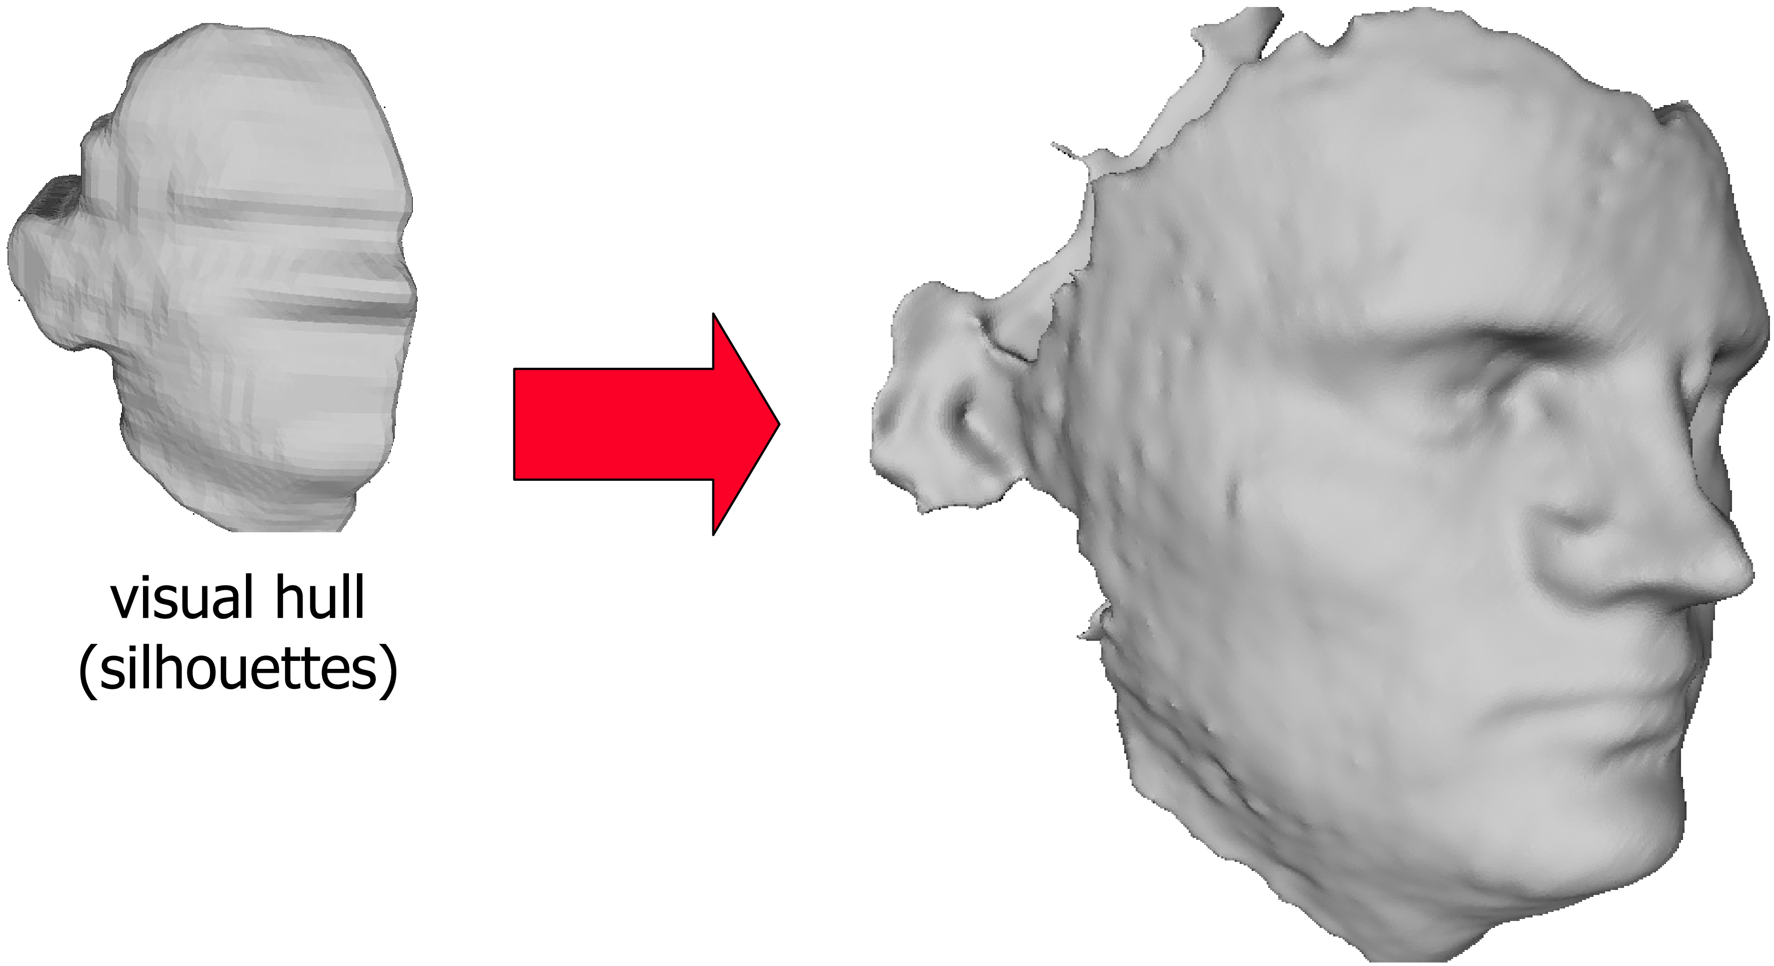

CVPR'05 slides from Vogiatzis, Torr, Cippola

Graph cuts for video textures

Graph-cuts video textures

(Kwatra, Schodl, Essa, Bobick 2003)

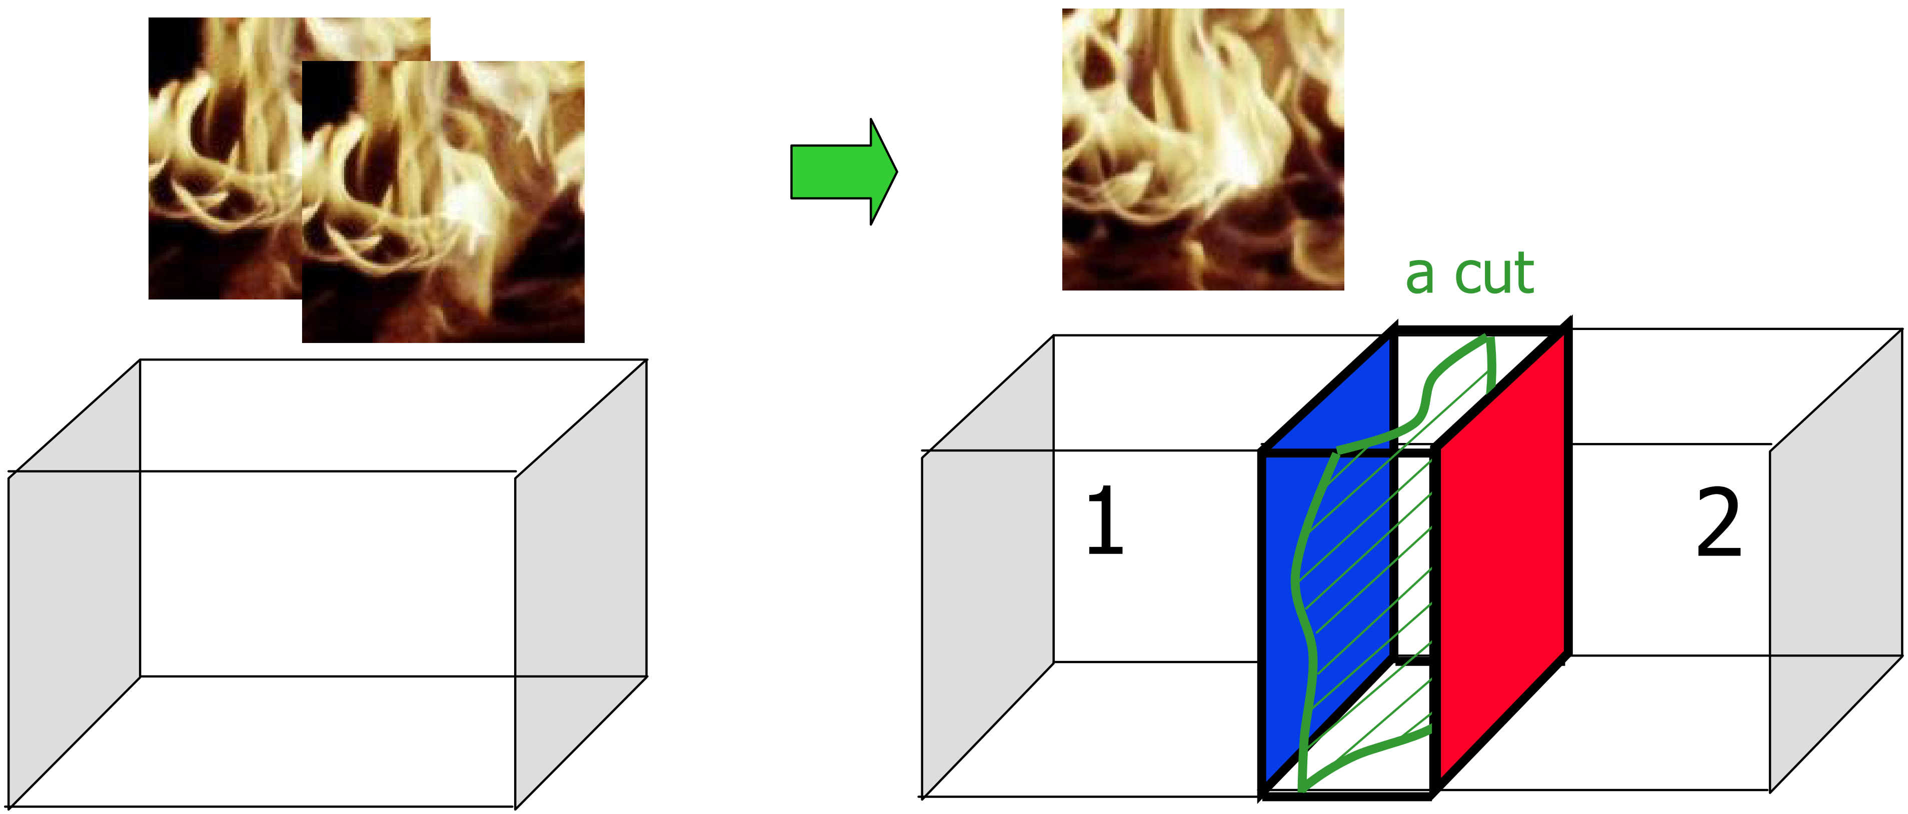

Short video clip

Long video clip

3D generalization of "**image-quilting**" (Efros & Freeman, 2001)

Graph cuts for video textures

Graph-cuts video textures

(Kwatra, Schodl, Essa, Bobick 2003)

original short clip

synthetic infinite texture

Cuts on directed graphs

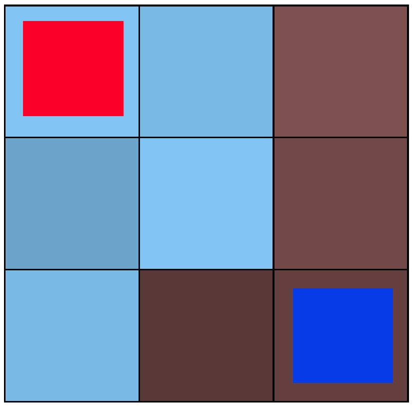

Flux of a vector field through hypersurface with orientation (A or B)

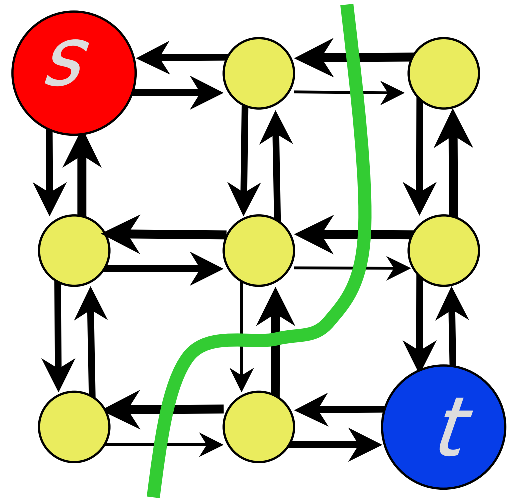

Cut on a directed graph

- Cost of a cut includes only edges from the source to the sink components
- Cut's cost (on a directed graph) changes if terminals are swapped

Swapping terminals *s* and *t* is similar to switching surface orientation

Cuts on directed graphs

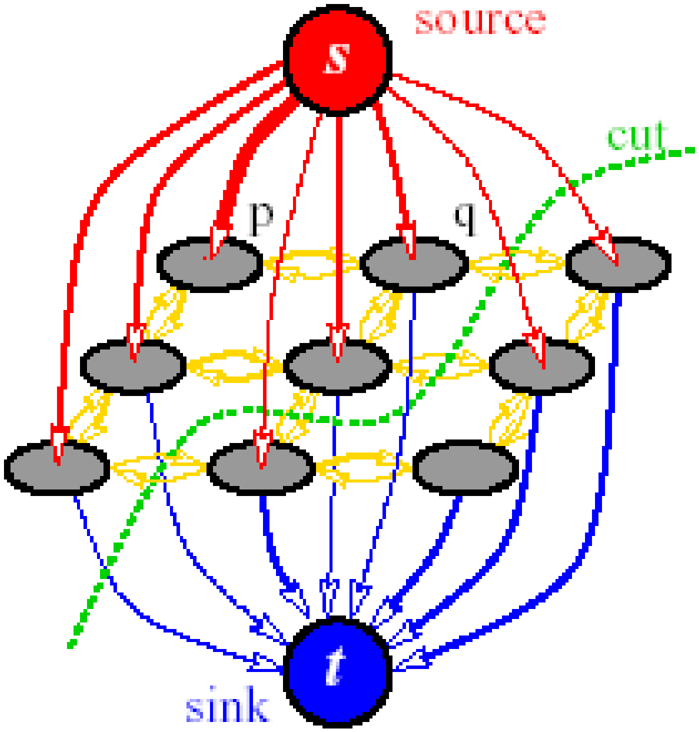

(a) directed graph

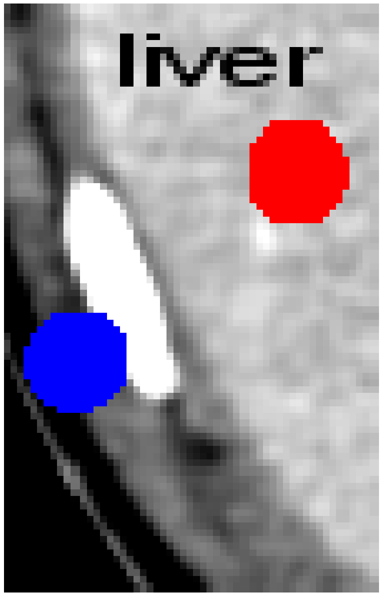

(b) image

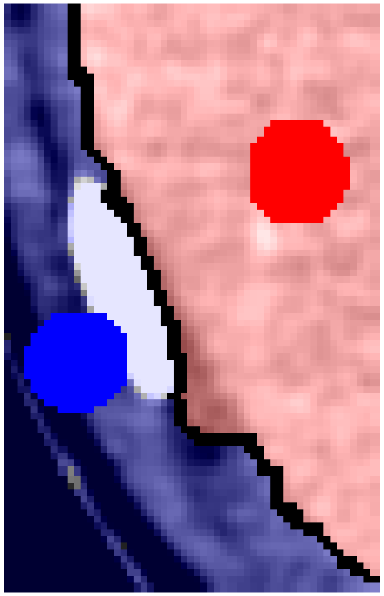

(c) undir. result

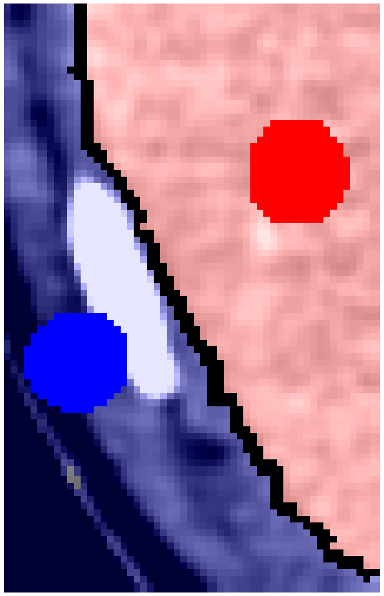

(d) dir. result

Segmentation of elongated structures

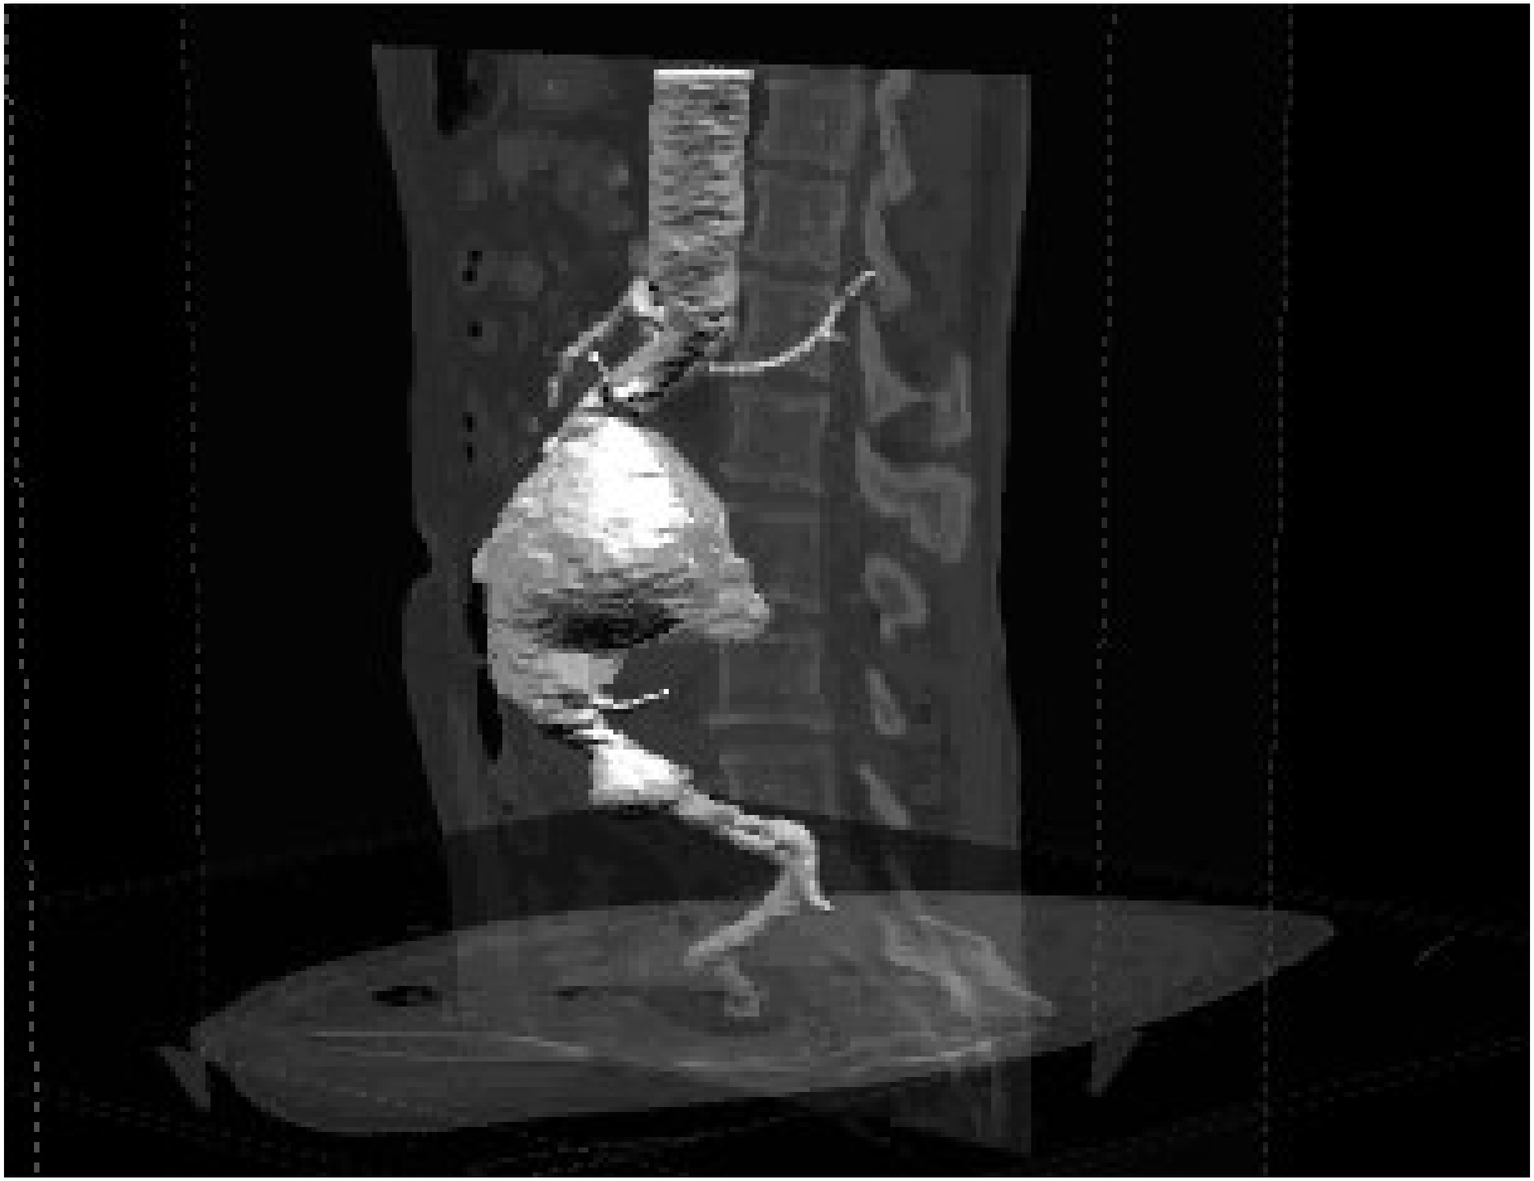

Simple “shape priors”

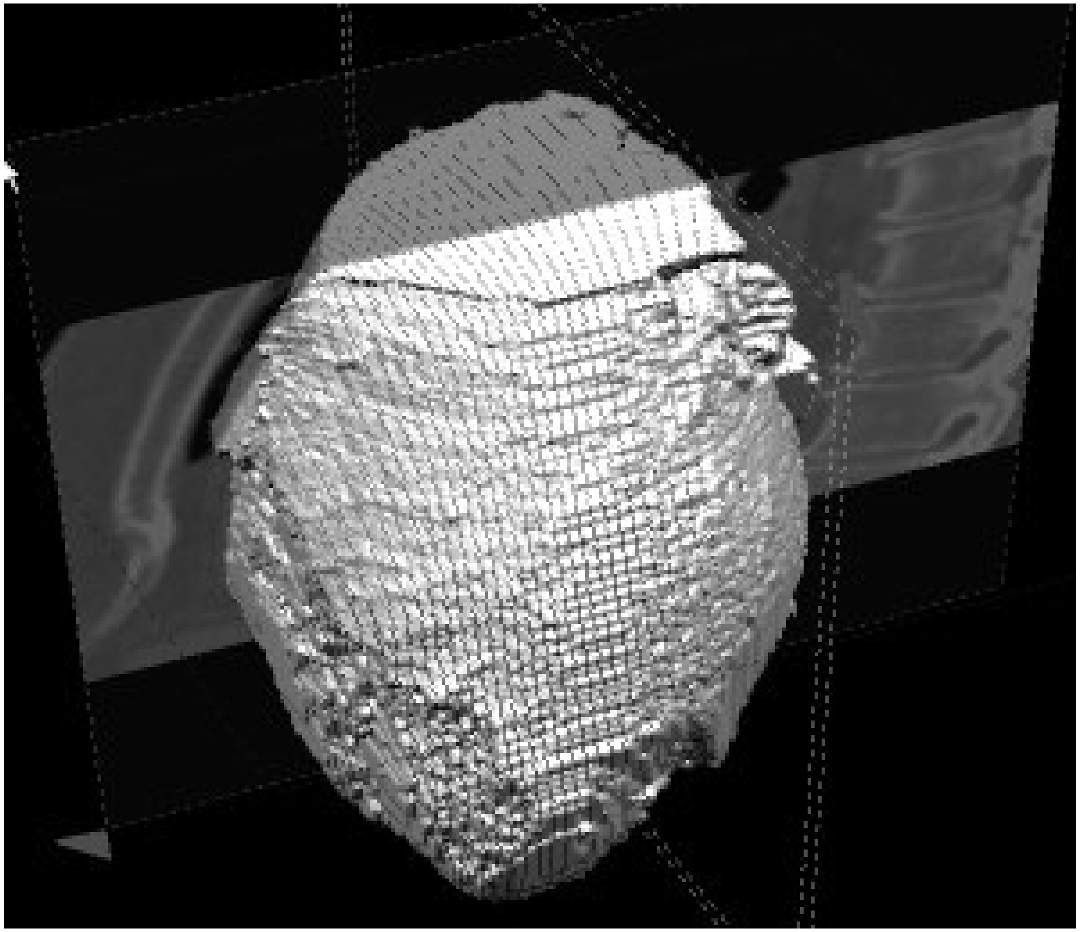

[Funka-Lea et al.’06]
“*blob*” prior

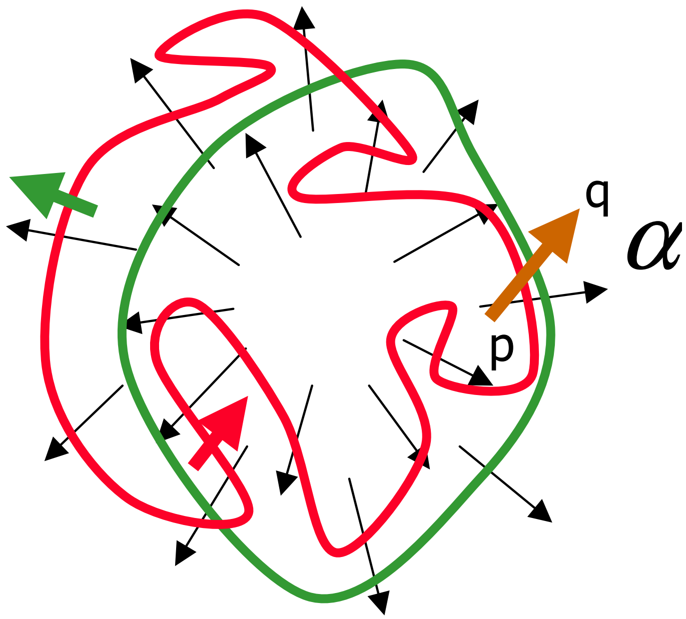

Extra penalty

$$w_{p->q} \propto f(\nabla I) + c \cdot (1 - \cos(\alpha))$$

Adding regional properties
(another segmentation example à la Boykov&Jolly'01)

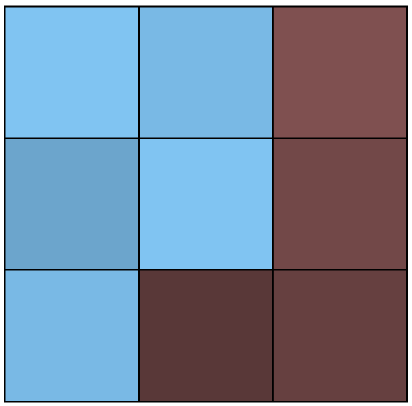

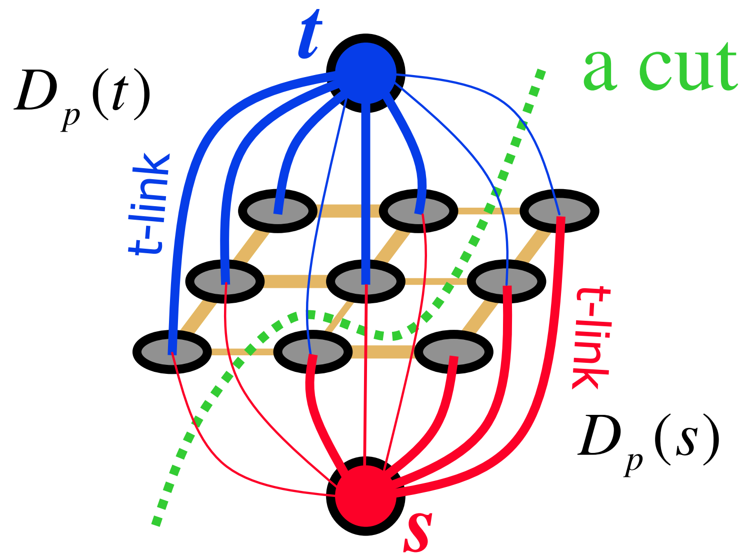

regional bias example

suppose I^s and I^t are given "expected" intensities of object and background

$$D_p(s) = const - | I_p - I^s |$$

$$D_p(t) = const - | I_p - I^t |$$

NOTE: hard constrains are not required, in general.

Adding regional properties
(another segmentation example à la Boykov&Jolly'01)

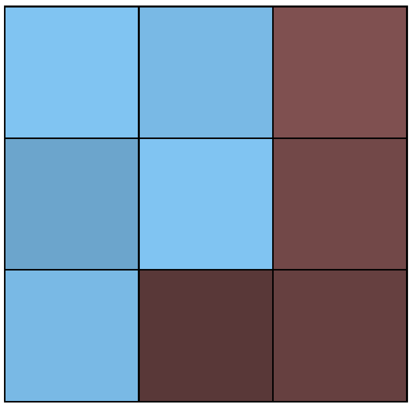

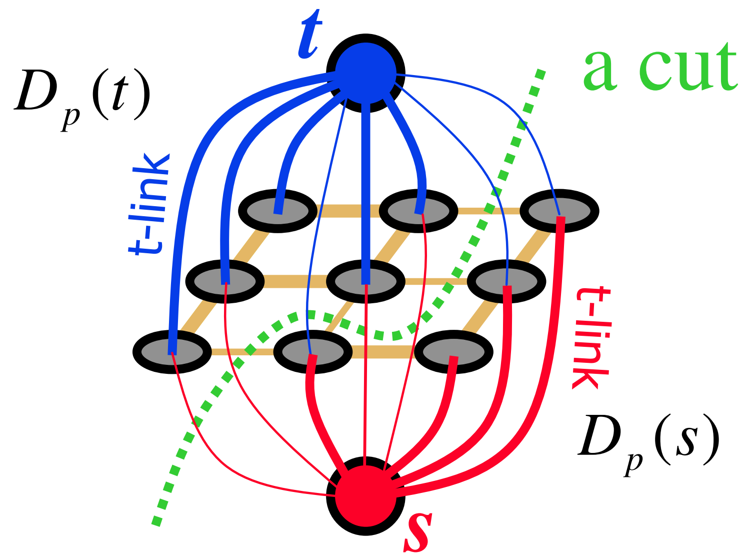

"expected" intensities of object and background I^s and I^t can be re-estimated

$$D_p(s) = const - | I_p - I^s |$$
$$D_p(t) = const - | I_p - I^t |$$

EM-style optimization of piece-vice constant *Mumford-Shah* model

Adding regional properties
(another example à la Boykov&Jolly, ICCV'01)

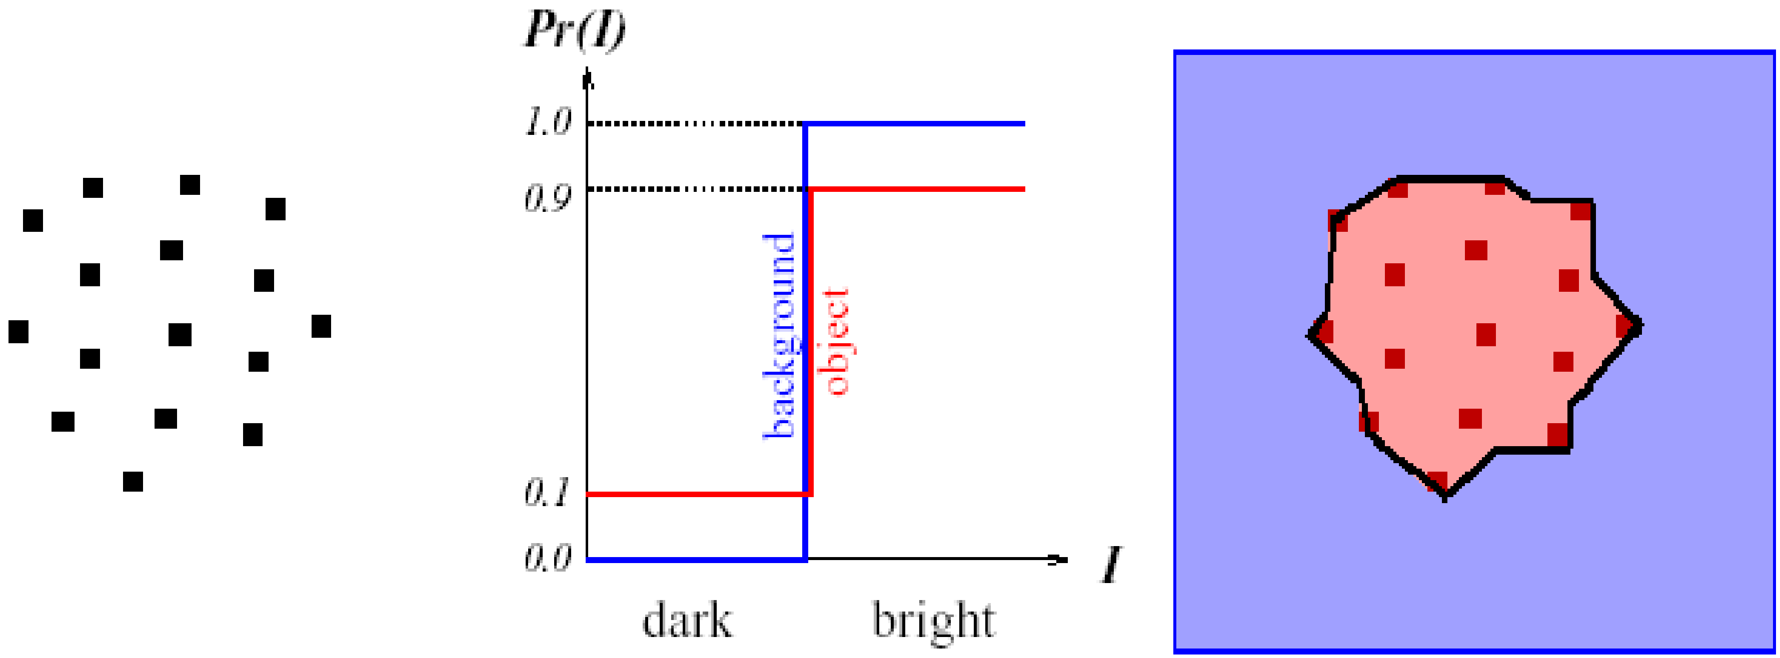

(a) Original image (b) Intensity histograms (c) Optimal segmentation

Adding regional properties
(another example à la Boykov&Jolly, ICCV'01)

More generally, regional bias can be based on any intensity models of object and background

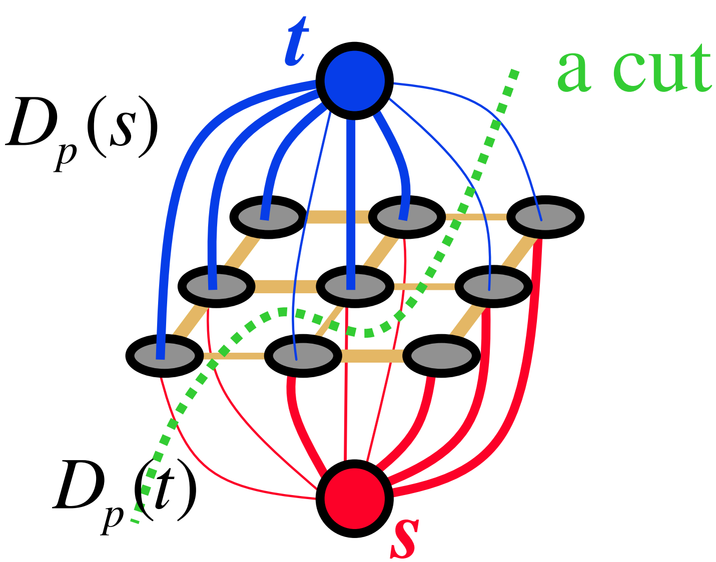

$$D_p(L_p) = -\ln \Pr(I_p \mid L_p)$$

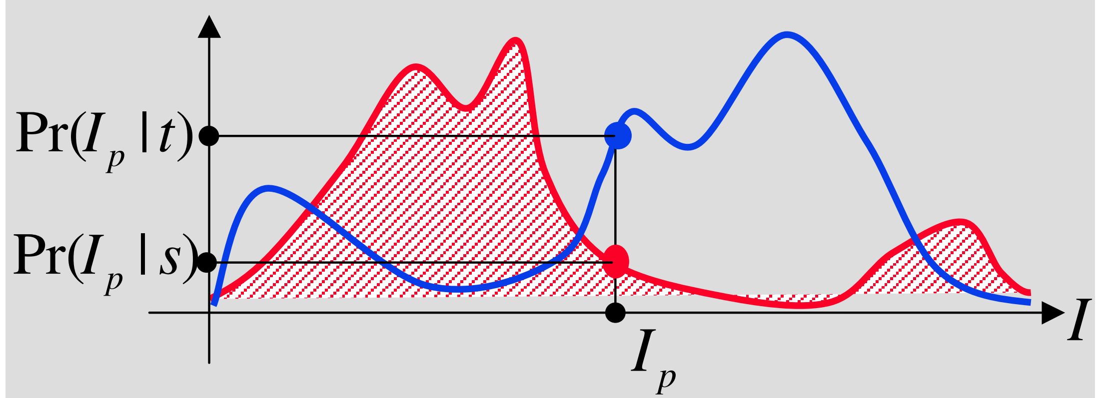

given object and background intensity histograms

Iterative learning of regional models

- GMMRF cuts (Blake et al., ECCV04)
- Grab-cut (Rother et al., SIGGRAPH 04)

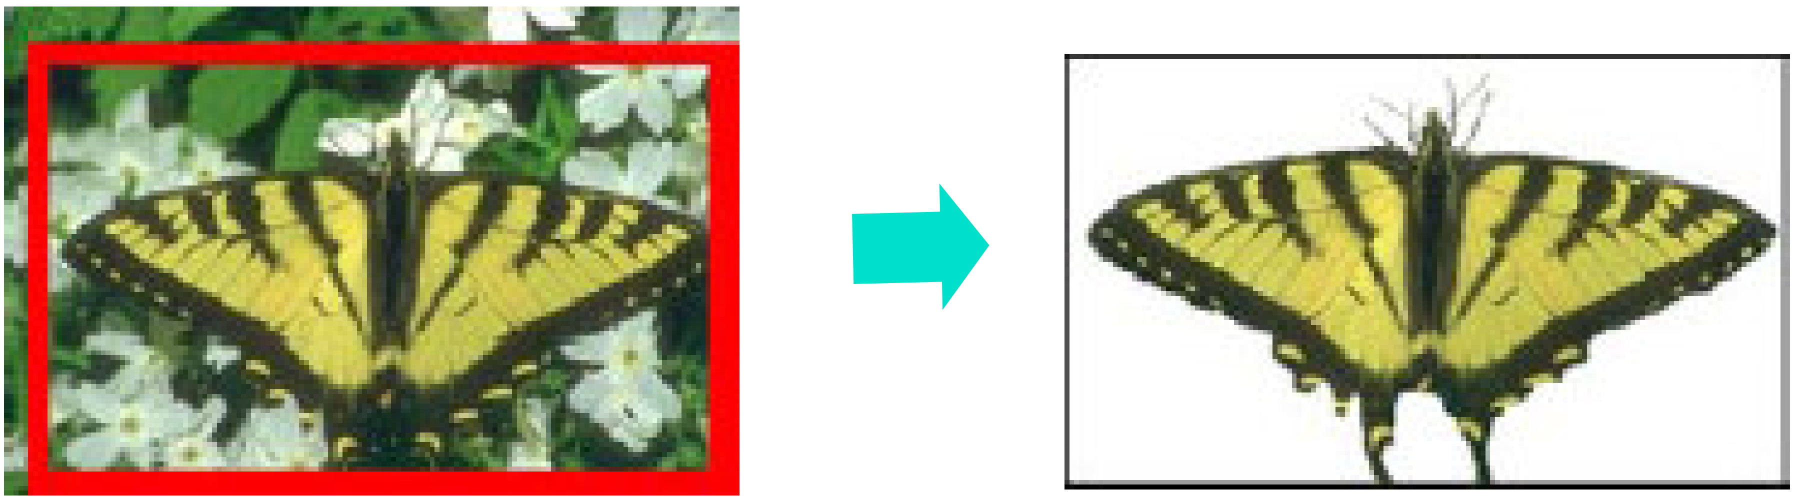

parametric regional model – Gaussian Mixture (GM)
designed to guarantee convergence

“Shrinking” bias

- Minimization of non-negative boundary costs (sum of non-negative n-links) gives regularization/smoothing of segmentation results
- Choosing n-link costs from local image gradients helps image-adaptive regularization

- May result in over-smoothing or “shrinking”
  - Typical for all surface regularization techniques
  - Graph cuts are no different from snakes or level-sets on that

"Shrinking" bias

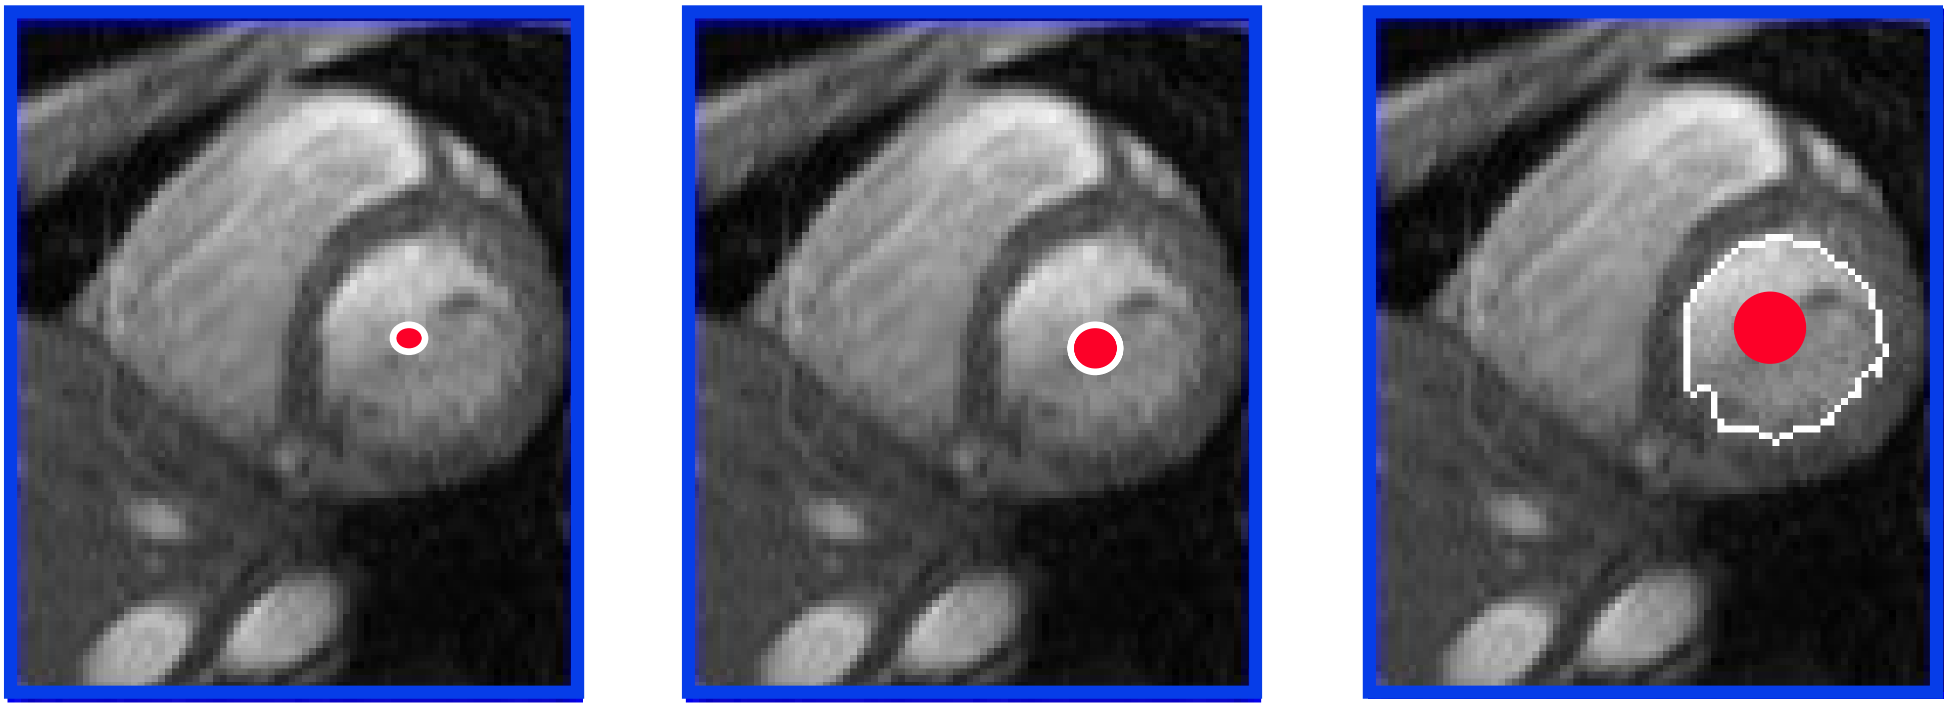

Optimal cut depending on the size of the hard constrained region

"Shrinking" bias

Image intensities on one scan line

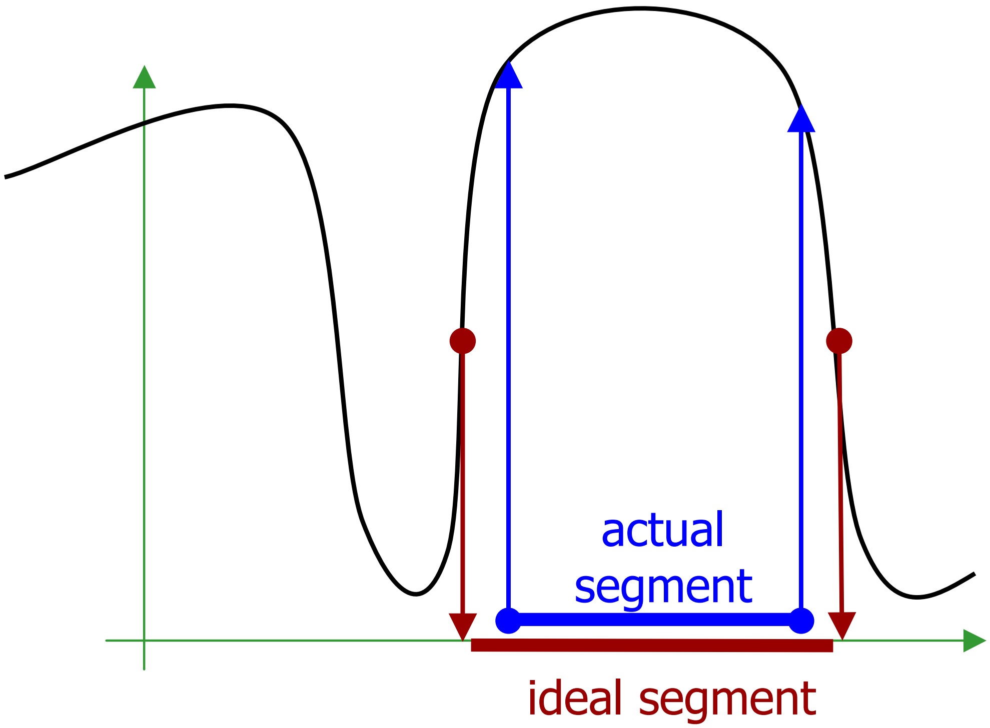

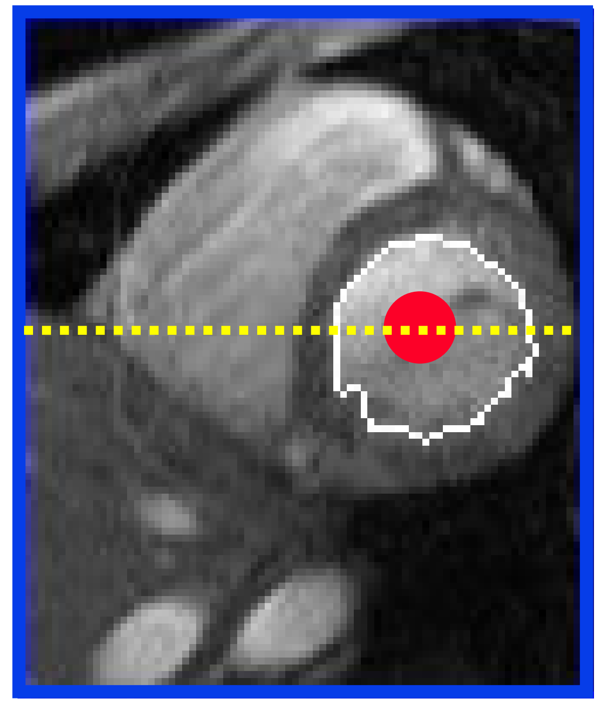

Shrinking (or under-segmentation)
is allowed by "intensity gradient ramp"

"Shrinking" bias

Surface regularization approach
to multi-view stereo

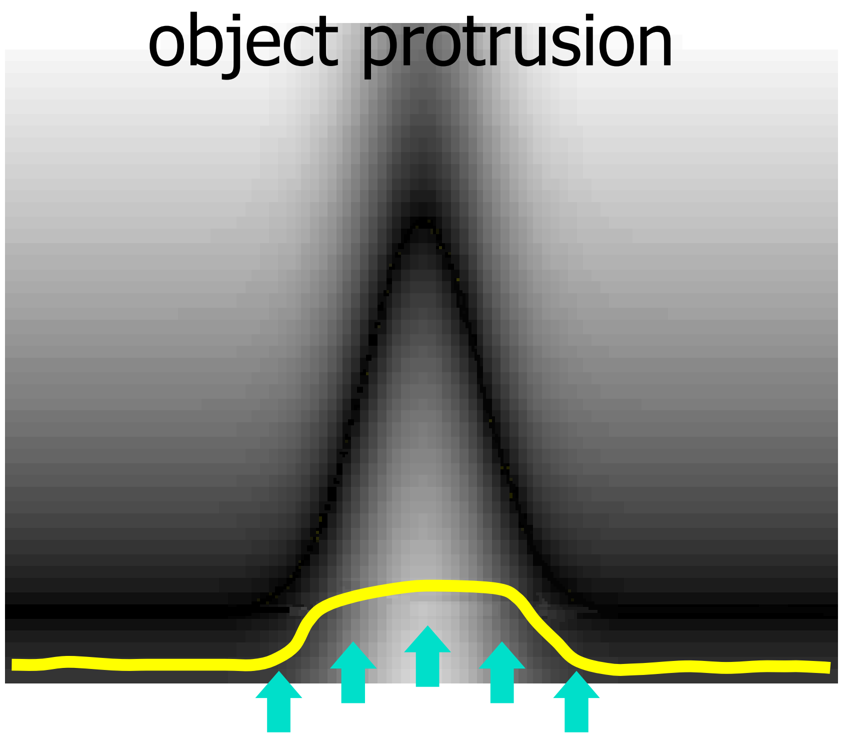

CVPR'05 slides from Vogiatzis, Torr, Cippola

Regional term can counter-act shrinking bias

All voxels in the center of the scene are connected to the object terminal

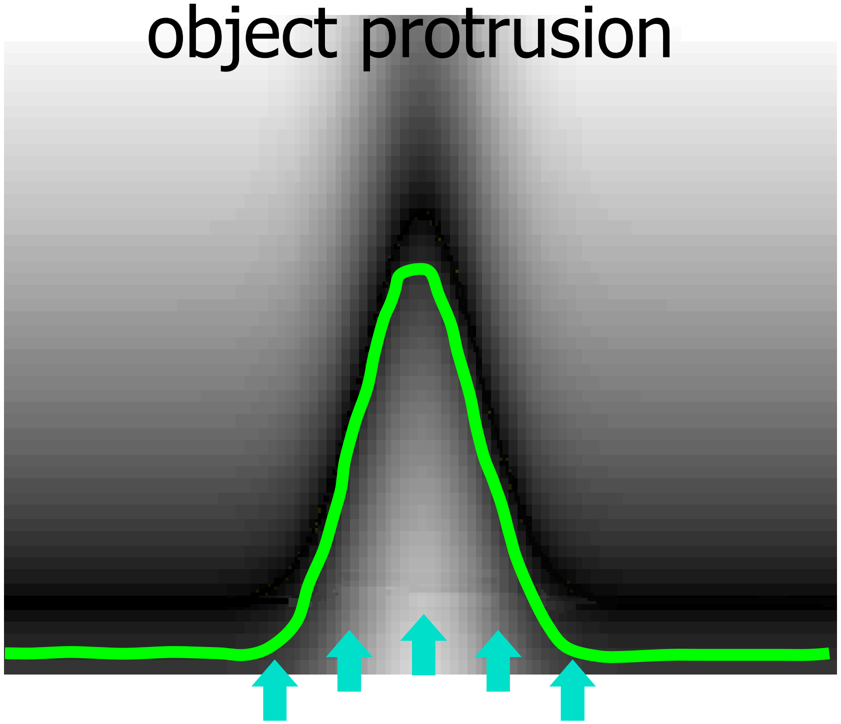

- “Ballooning” force
  - favouring bigger volumes that *fill* the visual hull

L.D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *PAMI*, 15(11):1131–1147, November 1993.

CVPR’05 slides from Vogiatzis, Torr, Cippola

"Shrinking" bias

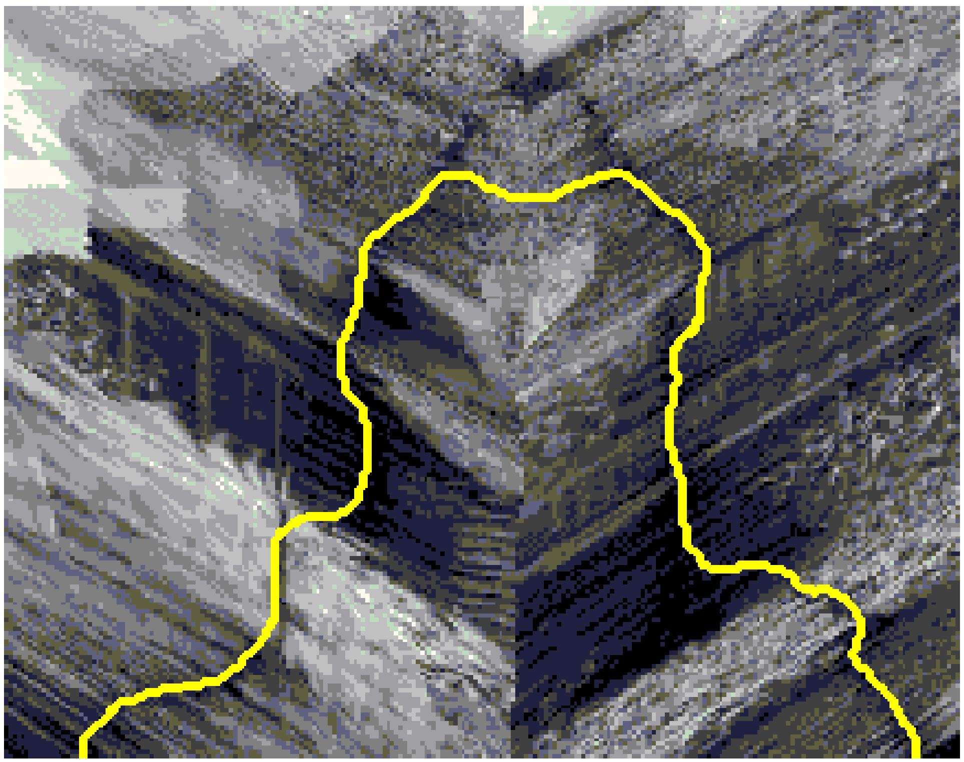

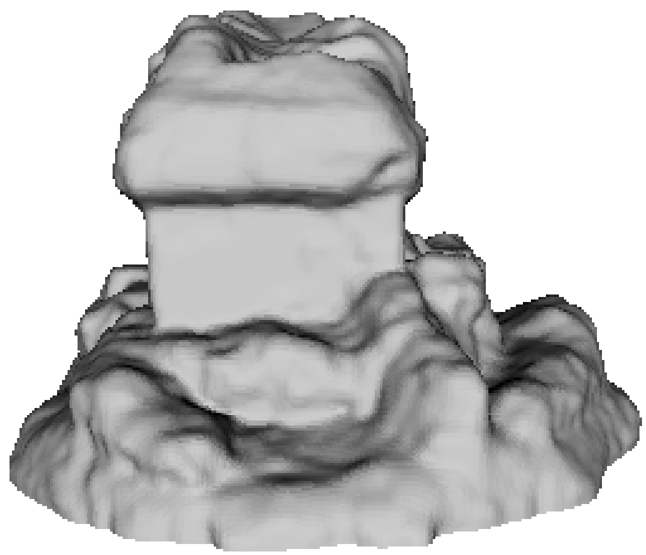

CVPR'05 slides from Vogiatzis, Torr, Cippola

Uniform ballooning




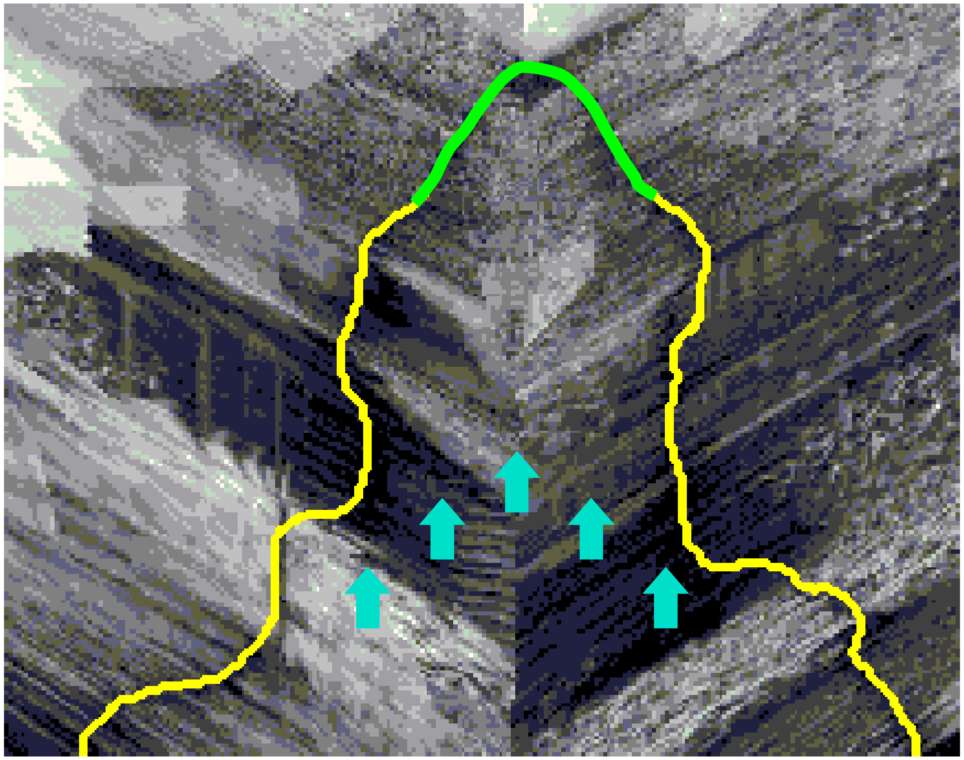
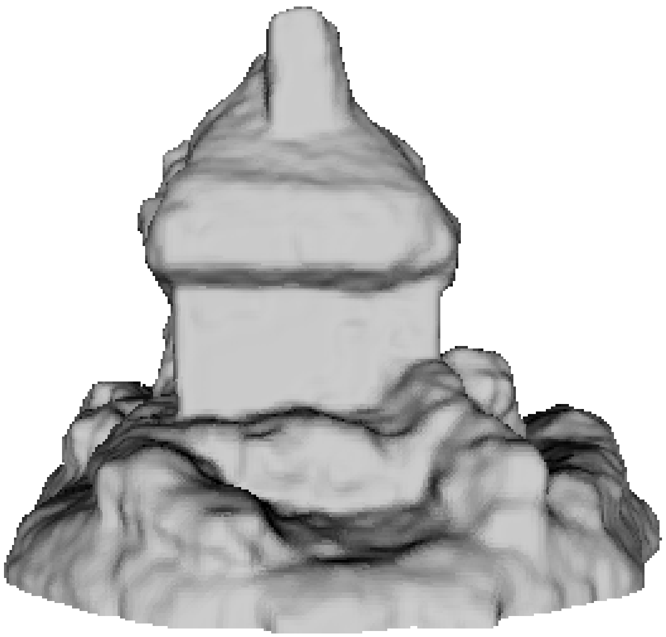

CVPR'05 slides from Vogiatzis, Torr, Cippola

Regional term based on Laplacian zero-crossings (*flux*)

Basic idea

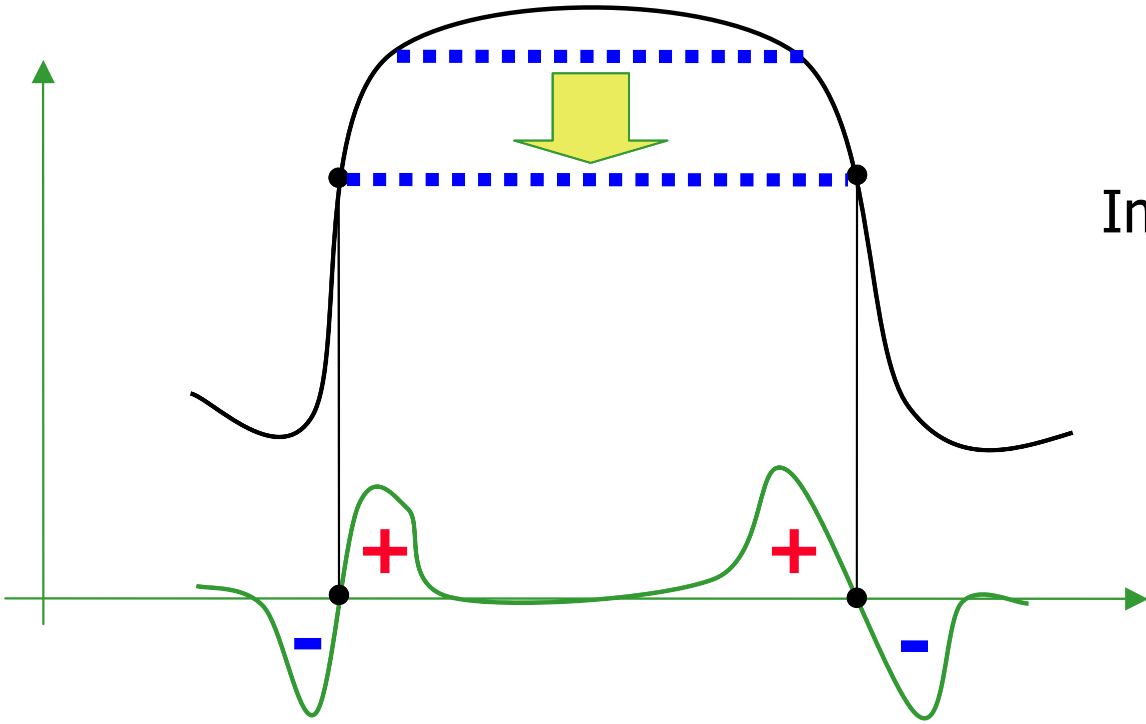

- Can be seen as intelligent ballooning

Vasilevsky and Sidiqqi, 2002

R. Kimmel and A. M. Bruckstein 2003

Integrating Laplacian zero-crossings into graph cuts (Kolmogorov&Boykov’05)

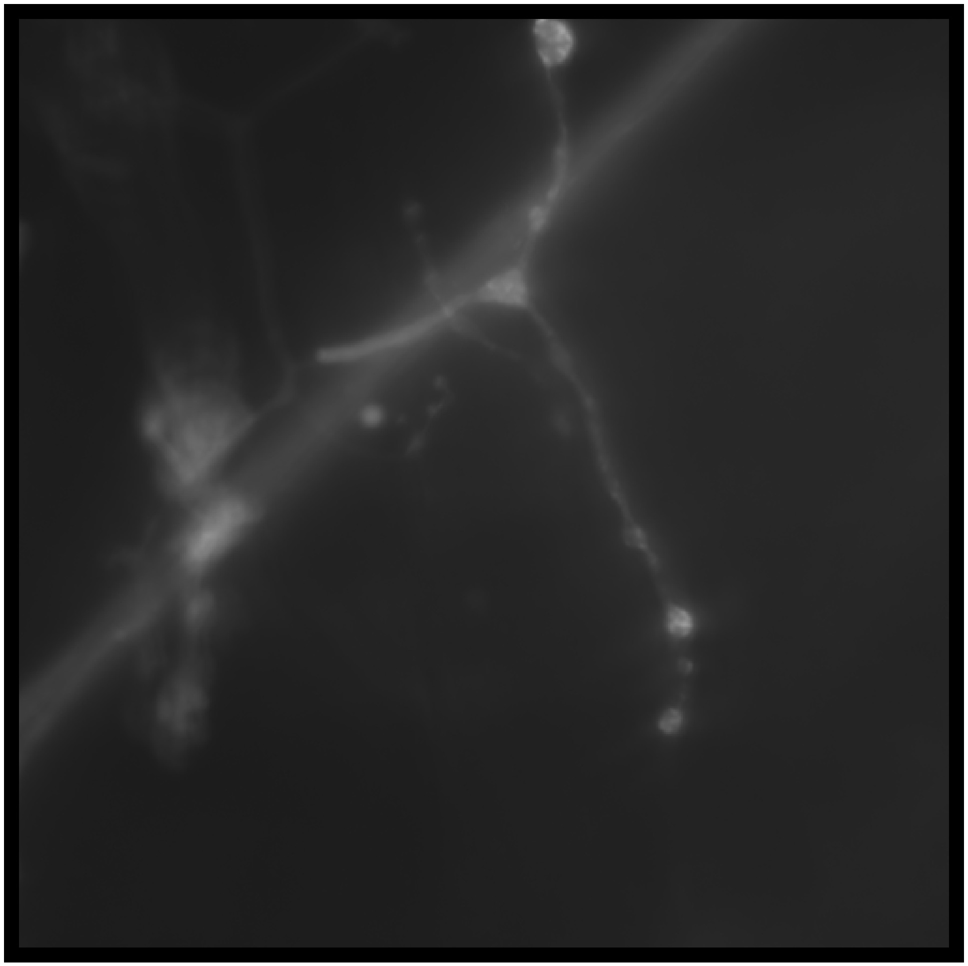

The image is courtesy of David Fleet
University of Toronto

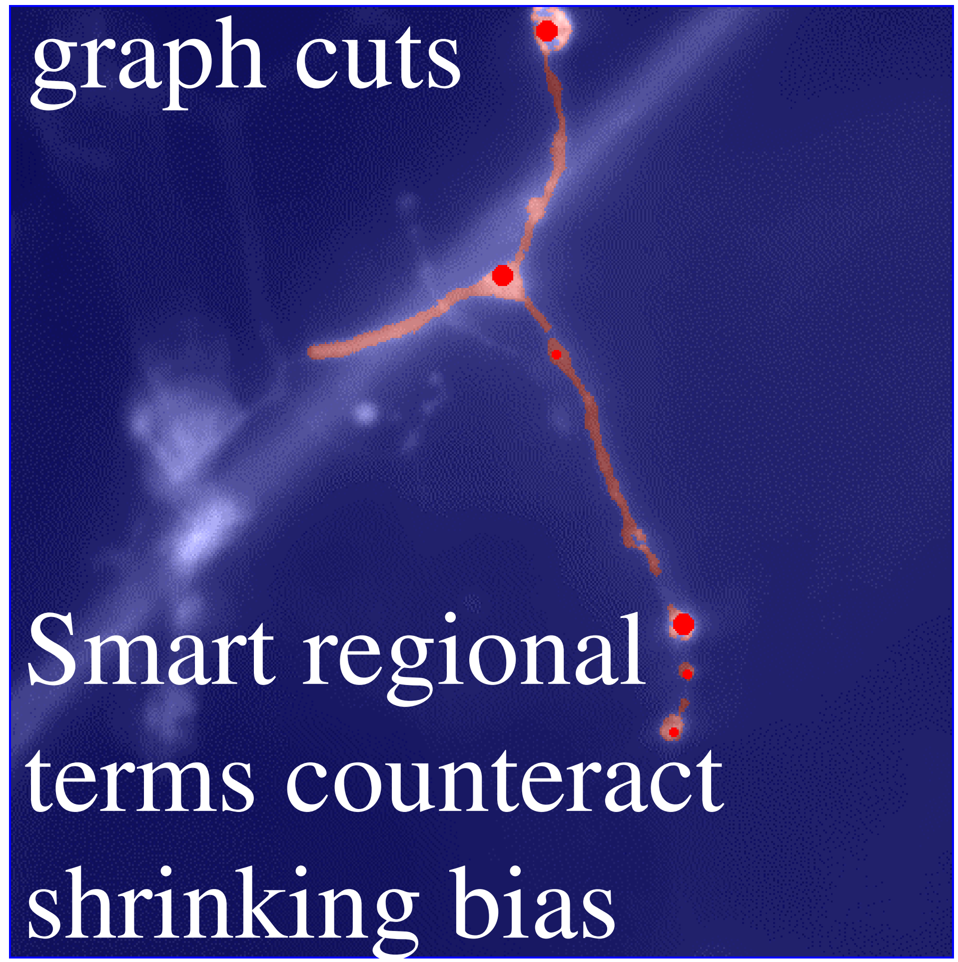

"Implicit" vs. "Explicit" graph cuts

- Most current graph cuts technique **implicitly** use surfaces represented via binary (interior/exterior) labeling of pixels

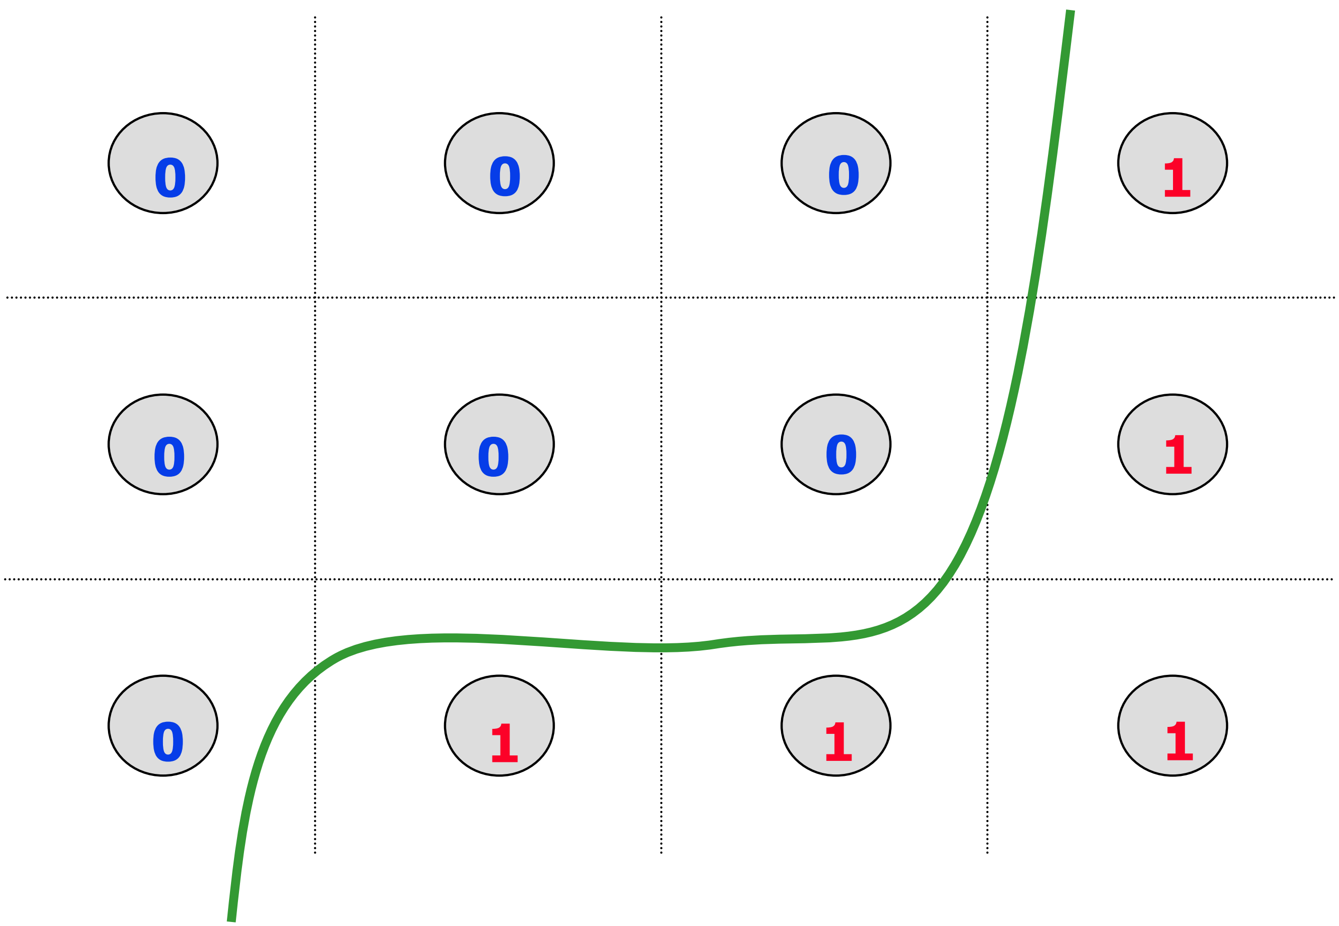

"Implicit" and "Explicit" graph cuts

- Except, a recent **explicit** surfaces representation method
  - Kirsanov and Gortler, 2004

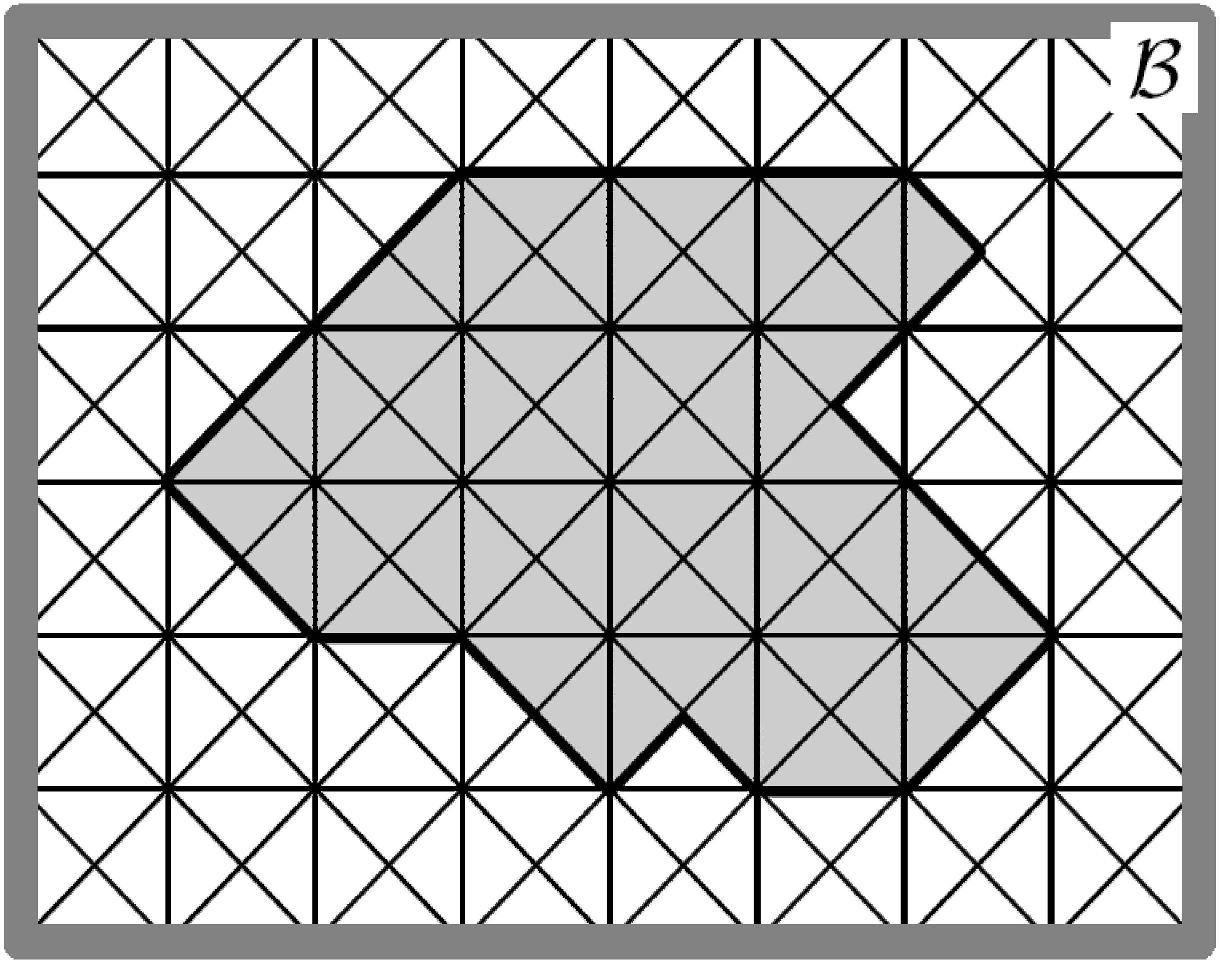

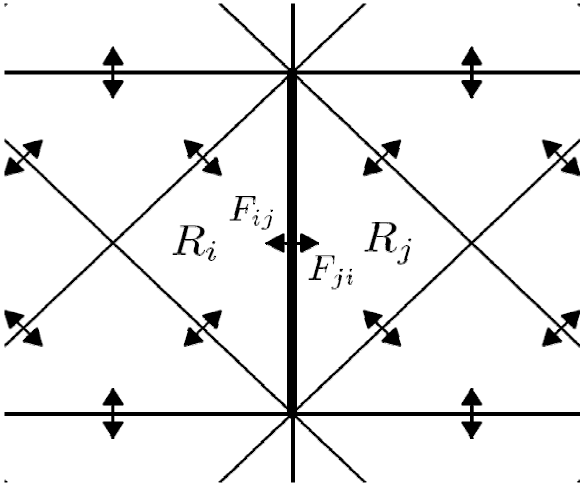

"Explicit" Graph Cuts

- for multi-view reconstruction
  - Lempitsky et al., ECCV 2006
- Explicit surface patches allow local estimation of visibility when computing globally optimal solution

local oriented visibility estimate

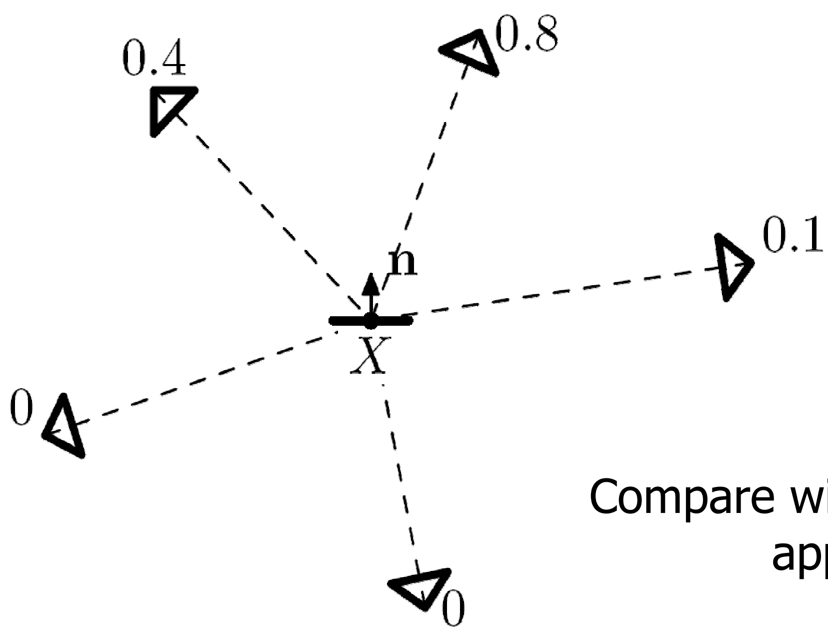

Compare with Vogiatzis et al., CVPR'05 approach for visibility

"Explicit" Graph Cuts

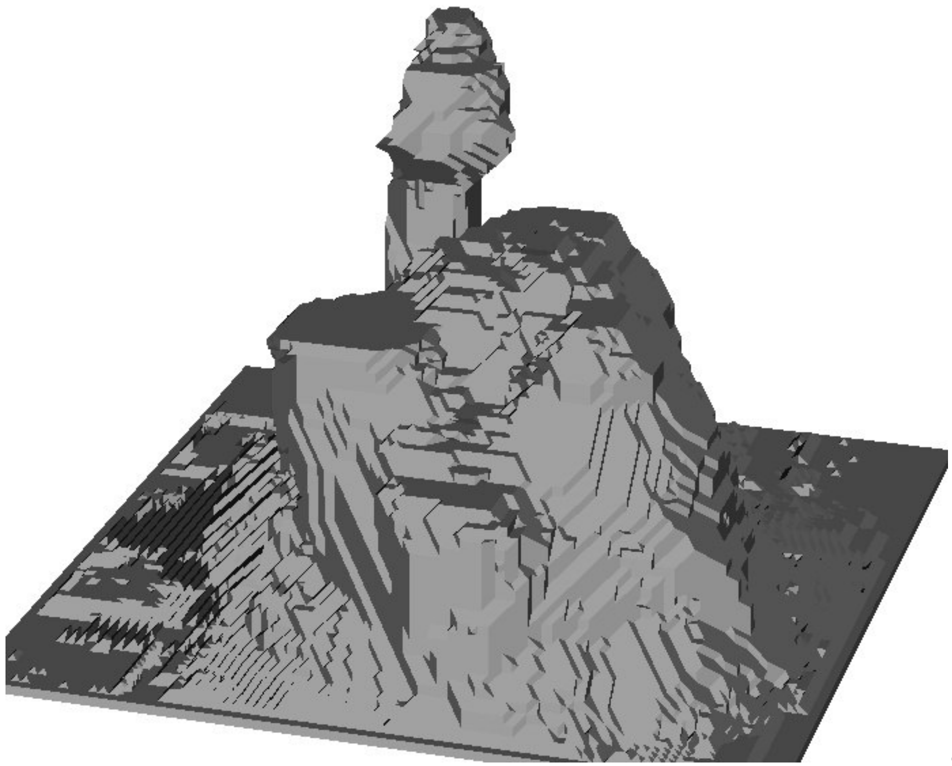

Regularization
+
uniform
ballooning

some details are
still over-smoothed

Lempitsky et al., ECCV 2006

"Explicit" Graph Cuts

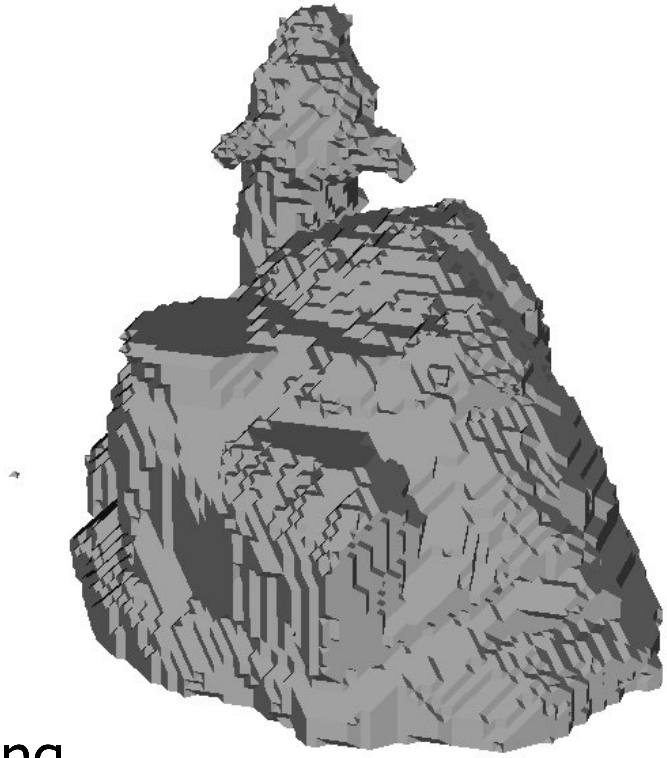

Regularization
+
intelligent
ballooning

Low noise and no shrinking

Boykov and Lempitsky, 2006

"Explicit" Graph Cuts

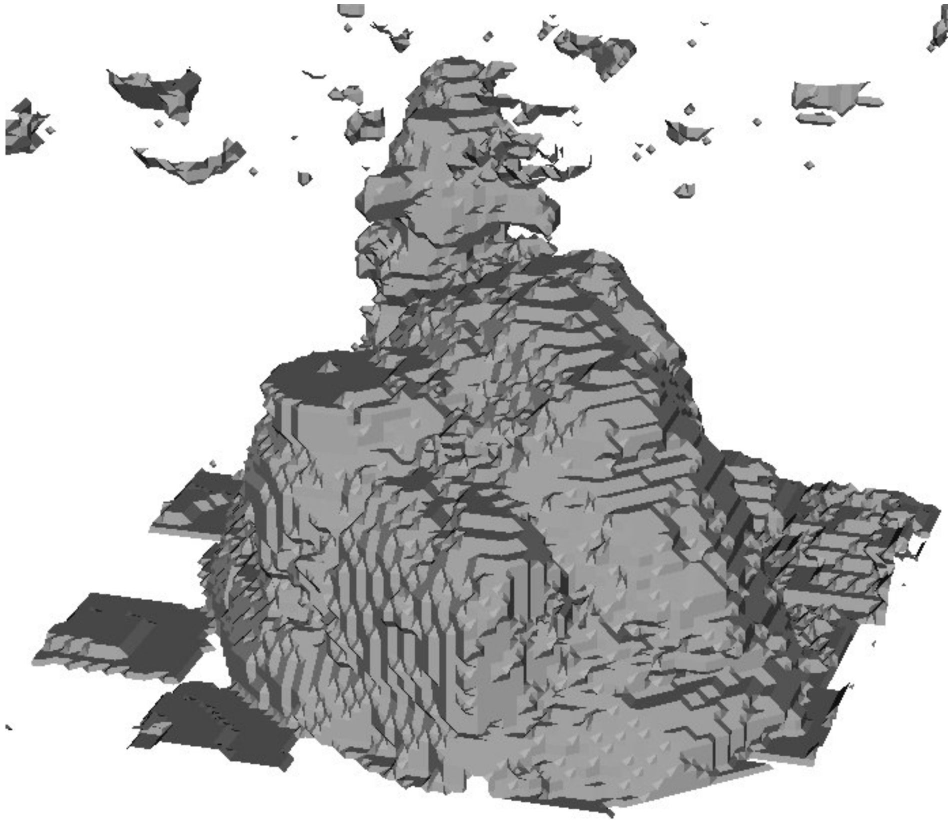

Space carving

Kutulakos and Seitz, 2000

Three ways to look at energy of graph cuts

I: Binary submodular energy

II: Approximating continuous surface functionals

III: Posterior energy (MAP-MRF)

Simple example of energy

Regional term Boundary term

$$E(L) = \sum_p -D_p(L_p) + \sum_{pq\in N} w_{pq} \cdot \delta(L_p \neq L_q)$$

t-links **n-links** $L_p \in \{s, t\}$

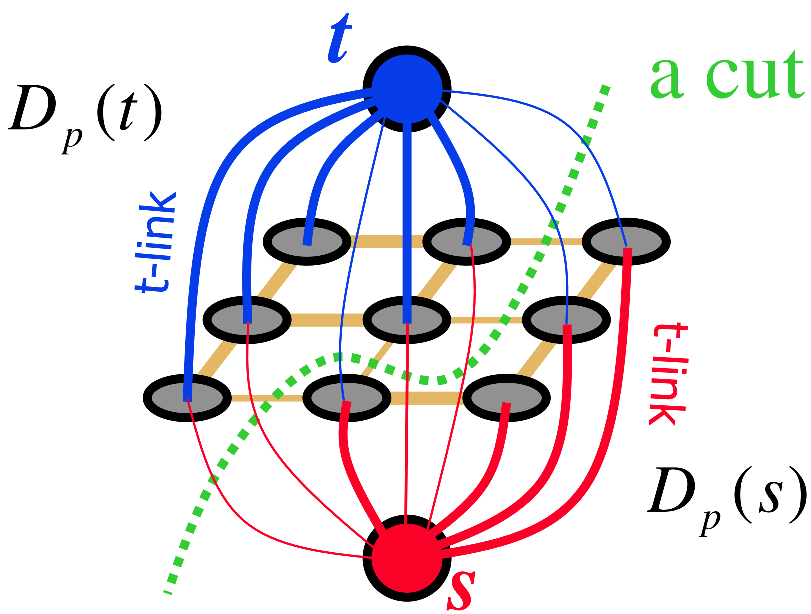

binary object segmentation

Graph cuts for minimization of submodular binary energies I

Regional term (t-links) — Boundary term (n-links)

$$E(L) = \sum_{p} E_p(L_p) + \sum_{pq \in N} E(L_p, L_q) \qquad L_p \in \{s,t\}$$

- Tight characterization of **binary** energies that can be globally minimized by *s-t* graph cuts is known
  [survey of Boros and Hummer, 2002, also Kolmogorov&Zabih 2002]

E(L) can be minimized by *s-t* graph cuts $\Leftrightarrow$ $E(s,s)+E(t,t) \le E(s,t)+E(t,s)$

submodularity condition for binary energies

- **Non-submodular cases** can be addressed with some optimality guarantees, e.g. ***QPBO*** algorithm reviewed in Boros and Hummer, 2002, nice slides by Kolmogorov at CVPR and Oxford-Brooks 2005 more in PART IV

Graph cuts for minimization of continuous surface functionals II

$$E(C) = \int_C g(\cdot)\, ds + \int_C \left\langle \vec{\mathbf{N}}, \vec{\mathbf{v}}_x \right\rangle ds + \int_{\Omega(C)} f(x)\, dp$$

Geometric length
any convex,
symmetric metric ***g***
e.g. Riemannian

Flux
any vector field **v**

Regional bias
any scalar function ***f***

- <u>Tight</u> characterization of energies of <u>**binary**</u> cuts *C* as functionals of continuous surfaces

[Boykov&Kolmogorov, ICCV 2003]

[Kolmogorov&Boykov, ICCV 2005]

more in PART III

Graph cuts for minimization of posterior energy III

- Greig at. al. [IJRSS, 1989]
  - Posterior energy (MRF, Ising model)

$$E(L) \;=\; \sum_{p} -\ln \Pr(D_p \mid L_p) \;+\; \sum_{pq\in N} V_{pq}(L_p, L_q)$$

Likelihood (data term)

Spatial prior (regularization)

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

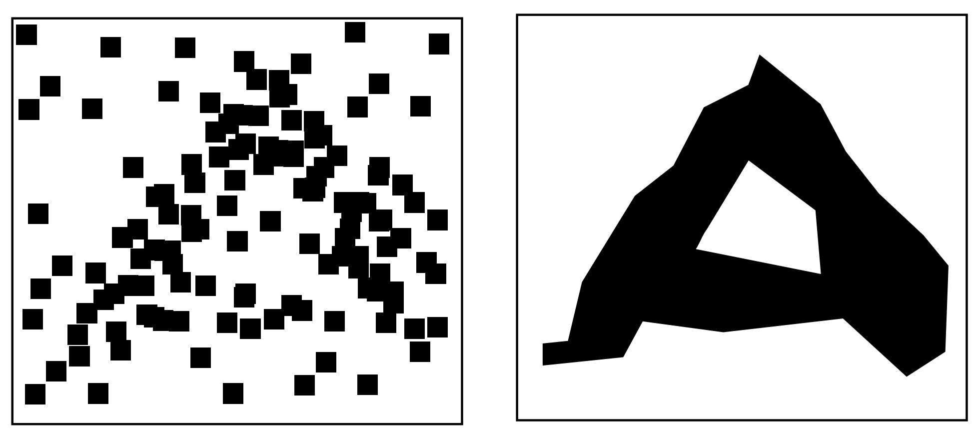

Example: binary image restoration

Graph cuts algorithms can minimize multi-label energies as well

Multi-scan-line stereo with *s-t* graph cuts (Roy&Cox'98)

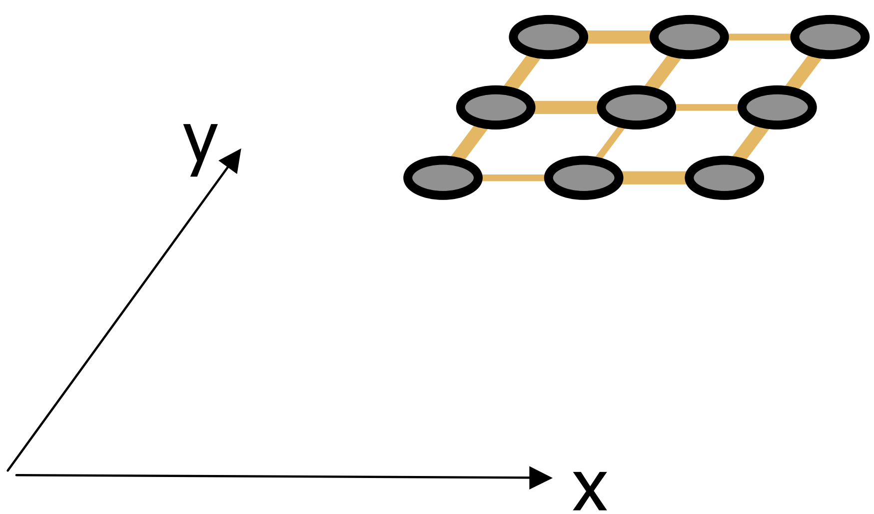

Multi-scan-line stereo with *s-t* graph cuts (Roy&Cox'98)

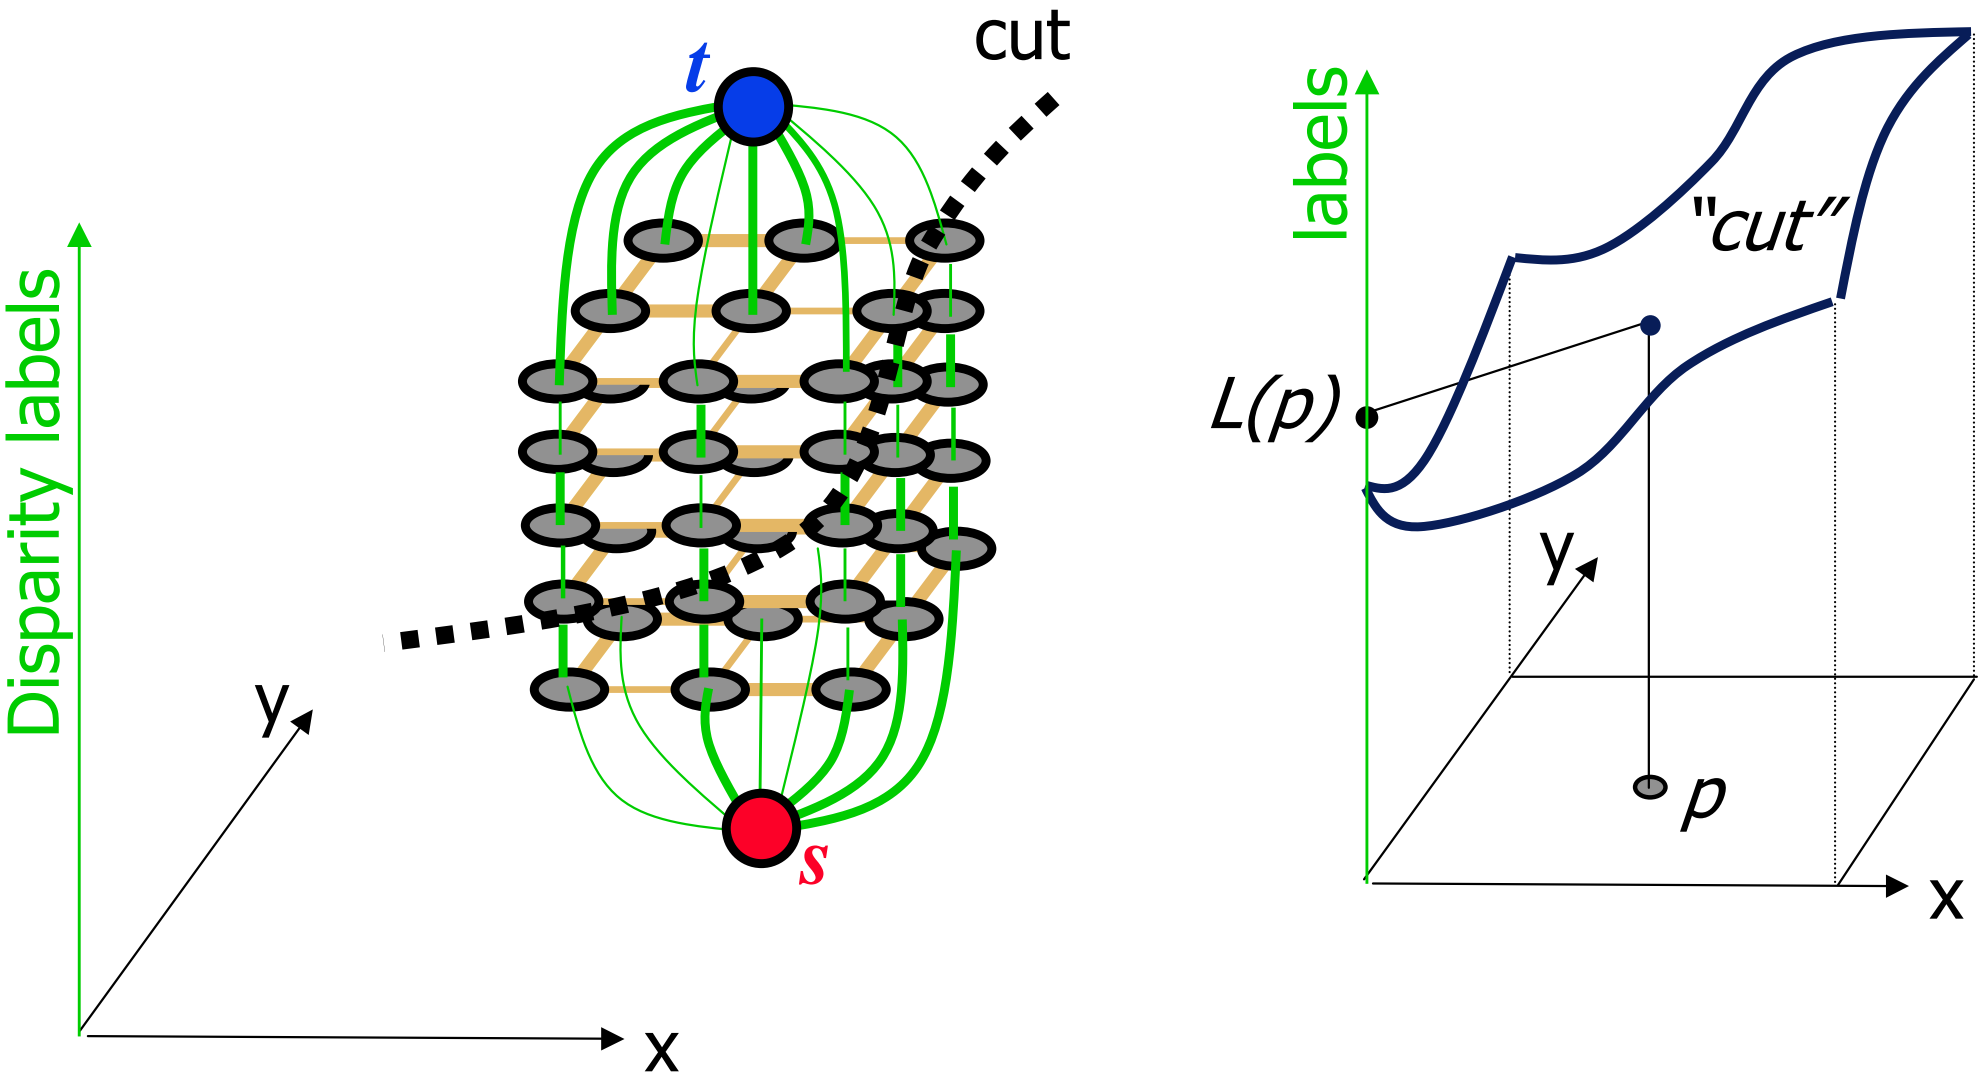

s-*t* graph-cuts for multi-label energy minimization

- Ishikawa 1998, 2000, 2003
- Modification of construction by Roy&Cox 1998

$$E(L) = \sum_{p} -D_p(L_p) + \sum_{pq \in N} V(L_p, L_q)$$

$$L_p \in R^1$$

Linear interactions

"Convex" interactions

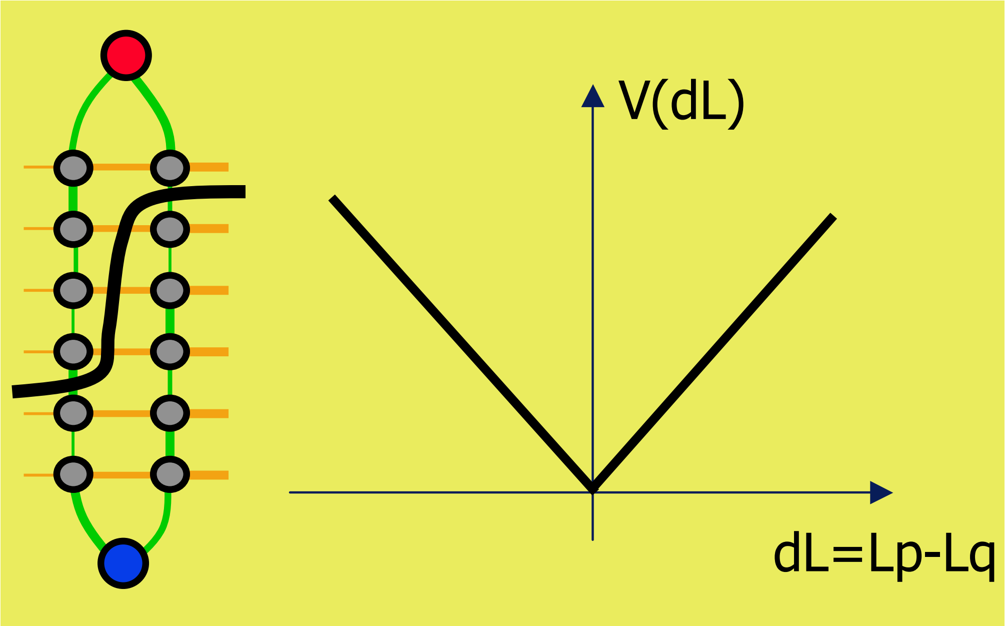

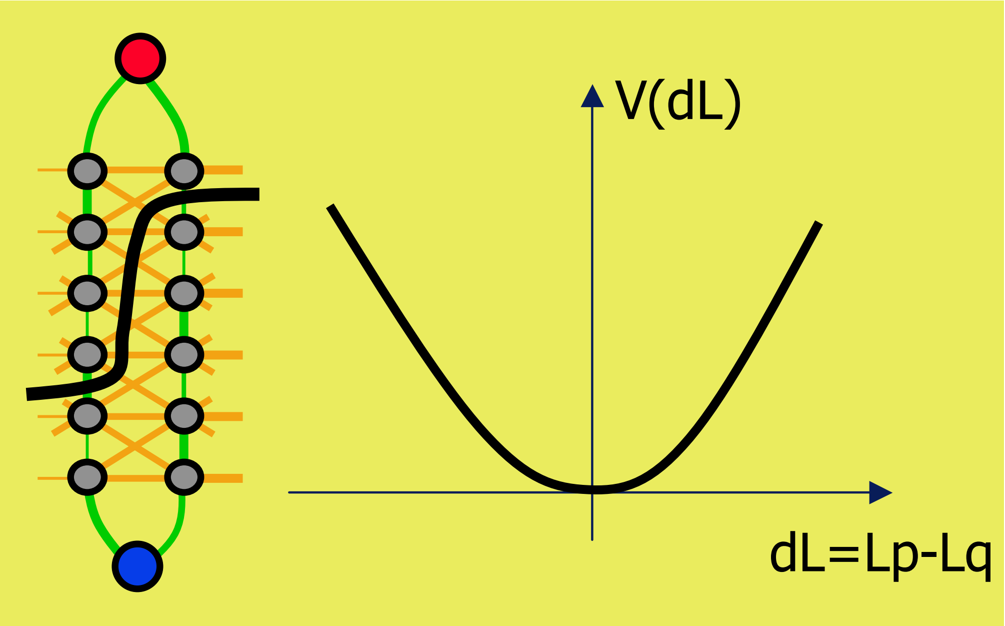

Pixel interactions *V*: “*convex*” vs. “*discontinuity-preserving*”

“Convex” Interactions *V*	Robust “discontinuity preserving” Interactions *V*
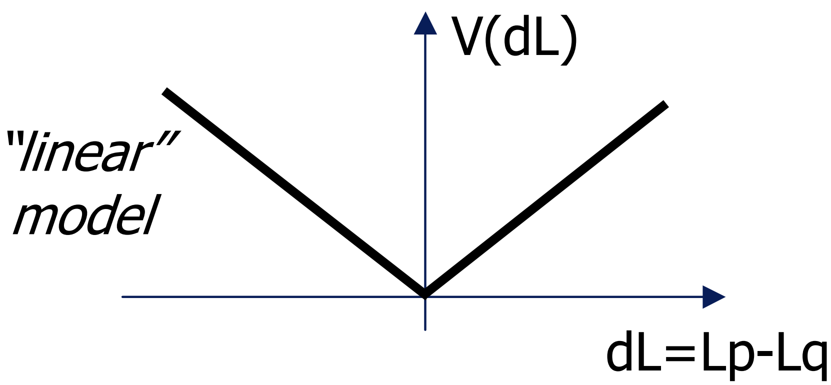	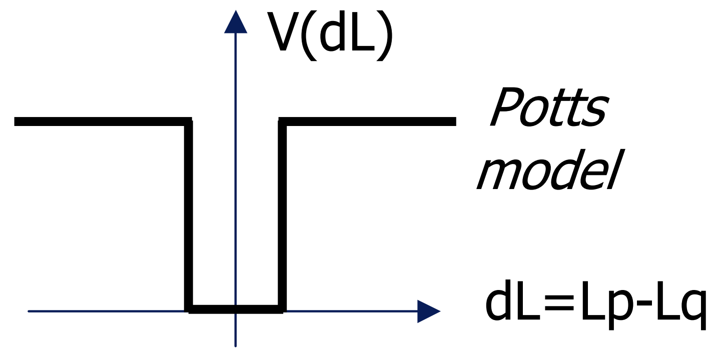
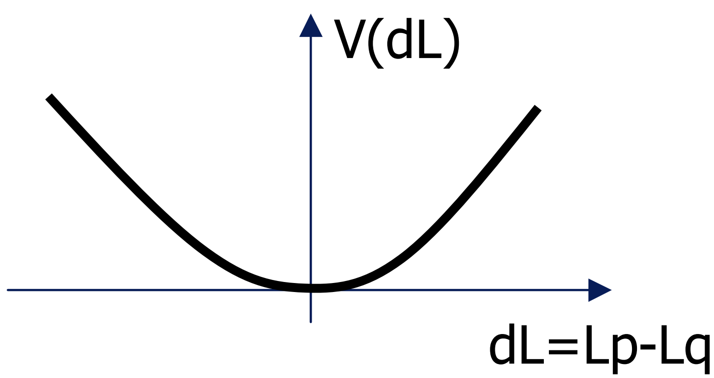	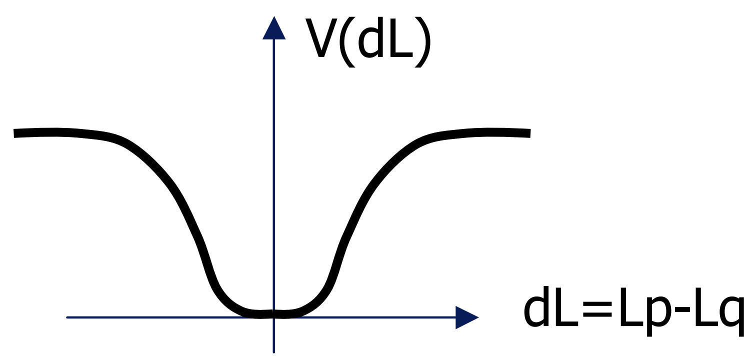

Pixel interactions: “*convex*” vs. “*discontinuity-preserving*”

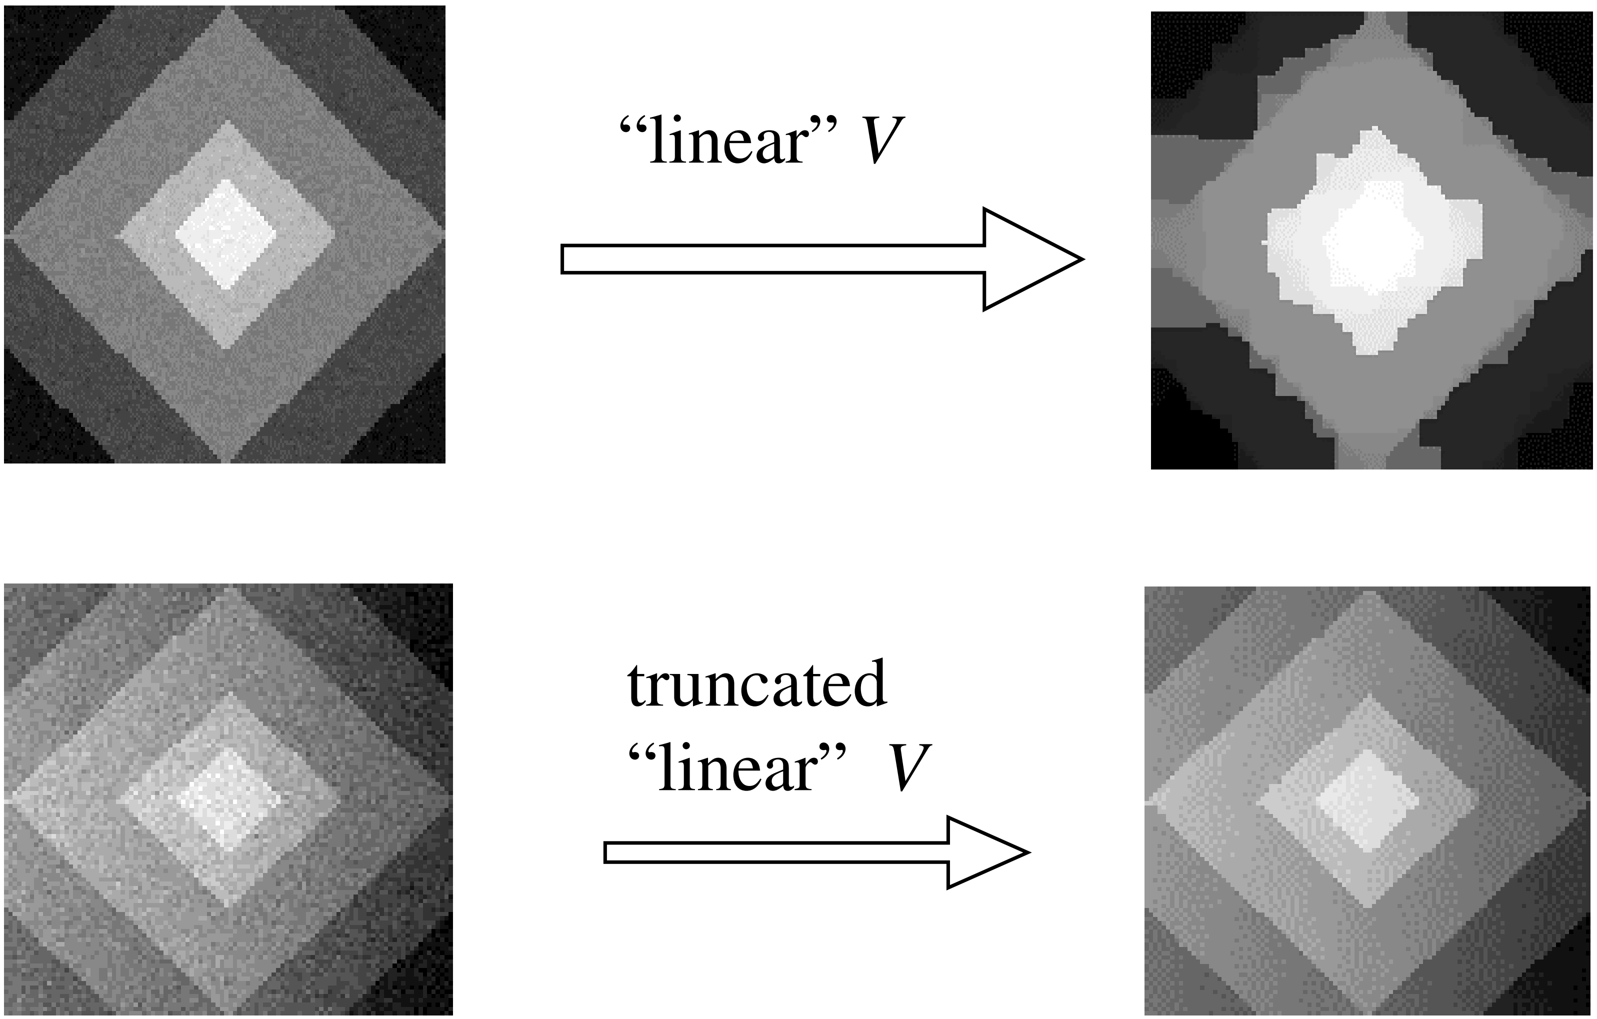

Robust interactions

- NP-hard problem (3 or more labels)
  - two labels can be solved via *s-t* cuts (Greig at. al., 1989)

- *a-expansion* approximation algorithm
  (Boykov, Veksler, Zabih 1998, 2001)
  - guaranteed approximation quality (Veksler, thesis 2001)
    - within a factor of 2 from the global minima (Potts model)
  - applies to a wide class of energies with robust interactions
    - Potts model (BVZ 1989)
    - "Metric" interactions (BVZ 2001)
    - "Submodular" interactions (e.g Boros and Hummer, 2002, KZ 2004)

a-expansion algorithm

1. Start with any initial solution
2. For each label "*a*" in any (e.g. random) order

    1. *Compute optimal a-expansion move (s-t graph cuts)*
    2. *Decline the move if there is no energy decrease*

3. *Stop when no expansion move would decrease energy*

a-expansion move

Basic idea: break multi-way cut computation into a **sequence of binary *s-t* cuts**

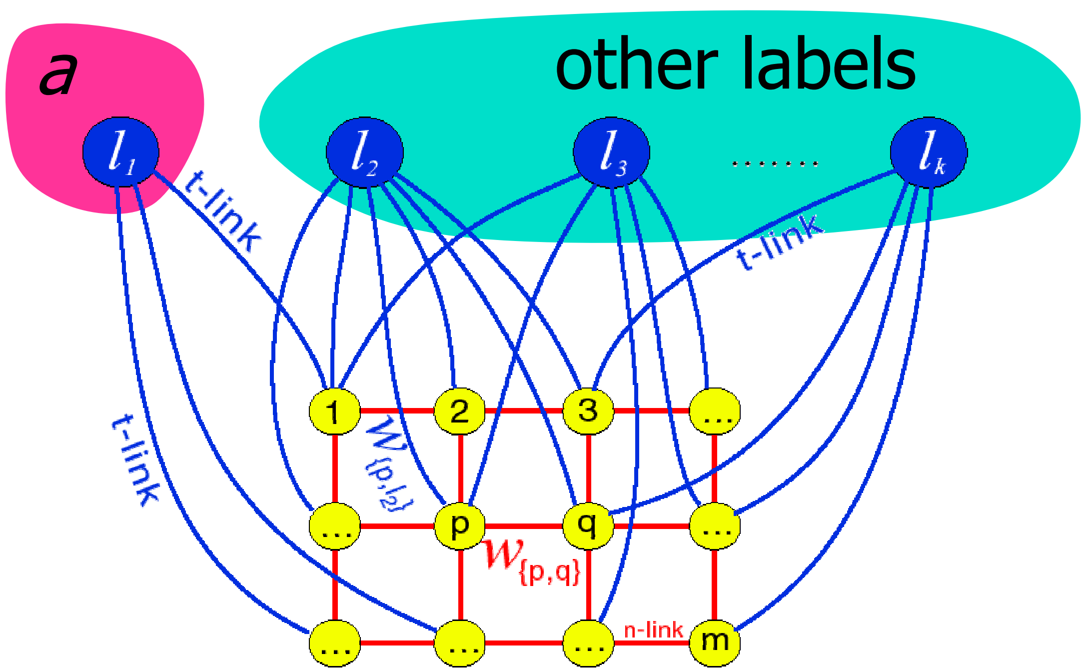

a-expansion moves

In each *a*-expansion a given label "*a*" grabs space from other labels

initial solution

- -expansion
- -expansion
- -expansion
- -expansion
- -expansion
- -expansion
- -expansion

For each move we choose expansion that gives the largest decrease in the energy: **binary optimization problem**

a-expansions: examples of *metric* interactions

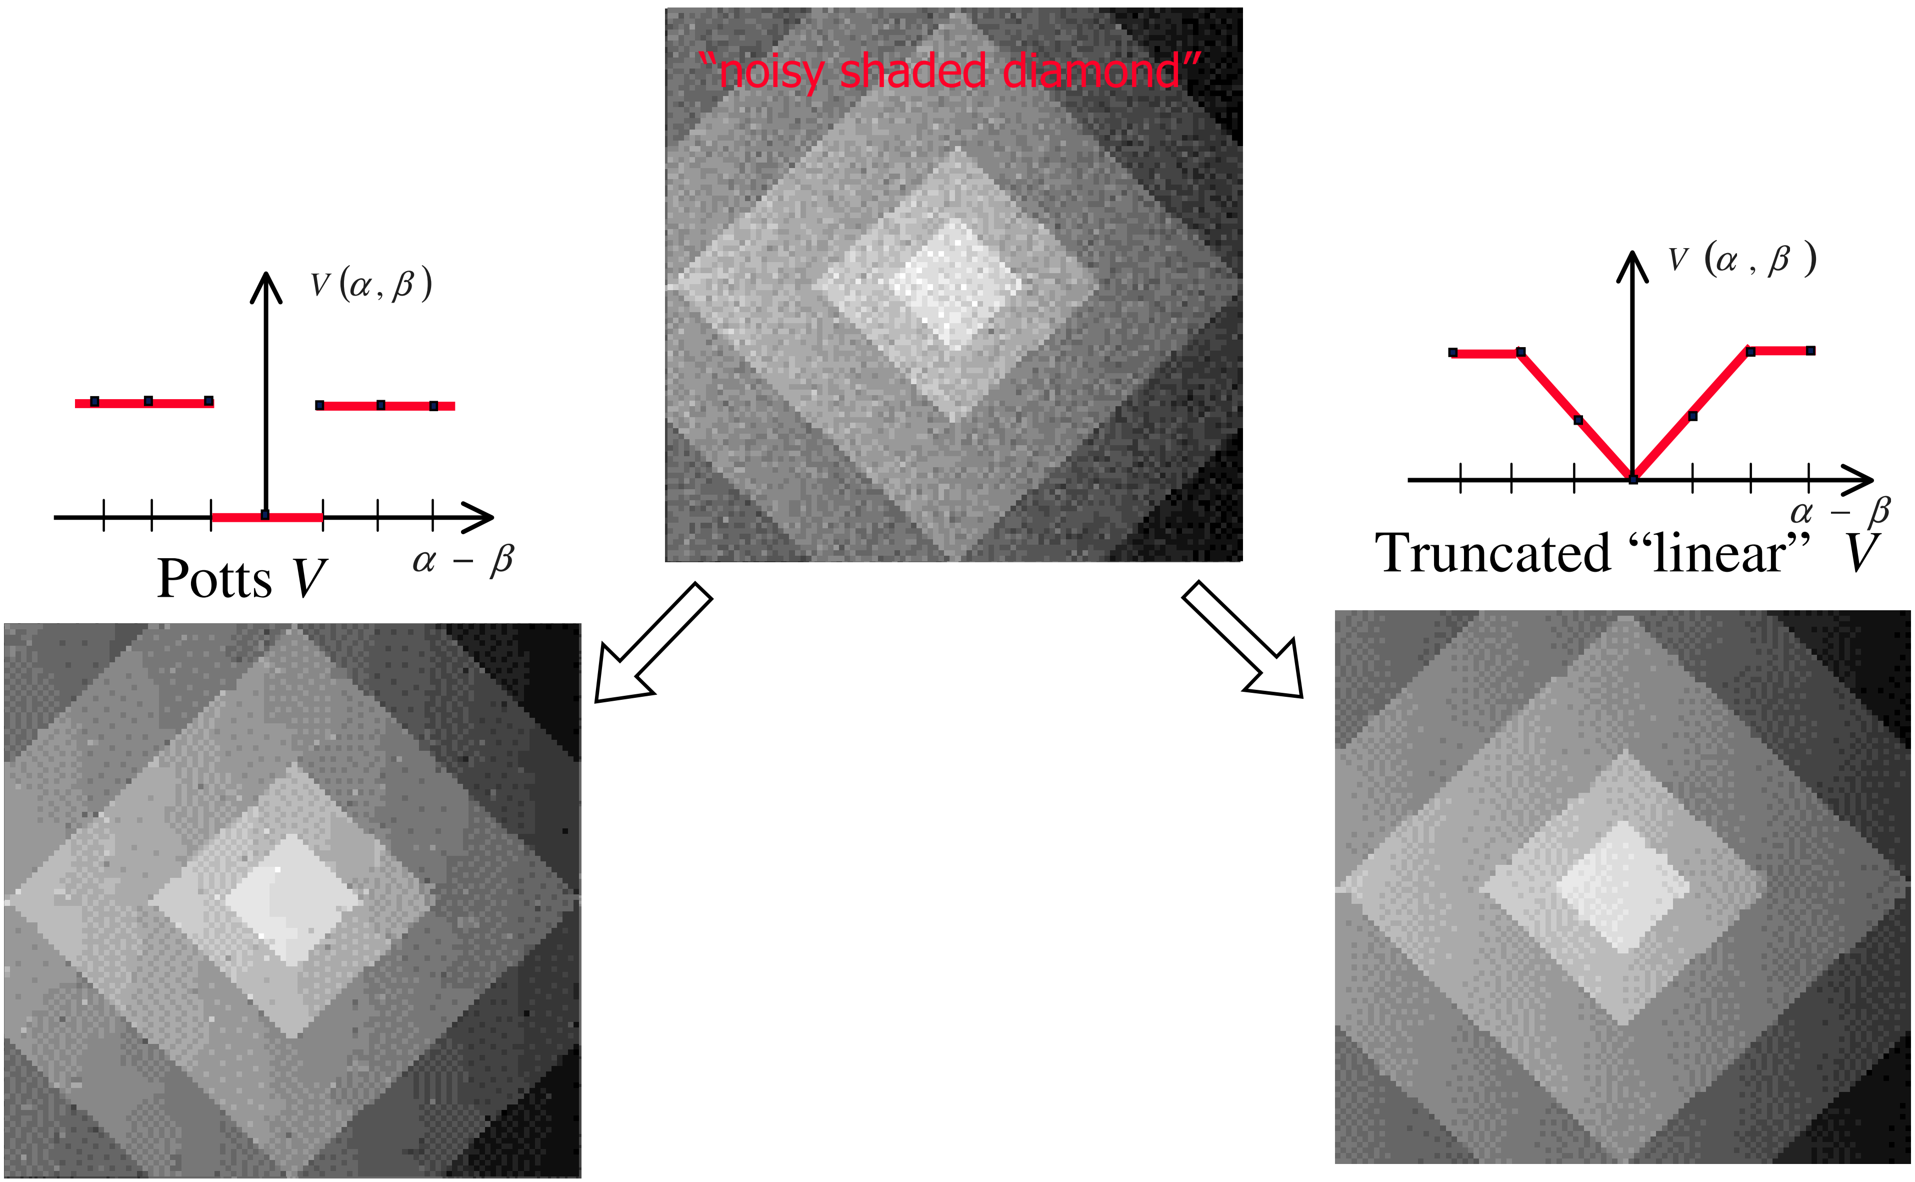

Multi-way graph cuts

Multi-object Extraction

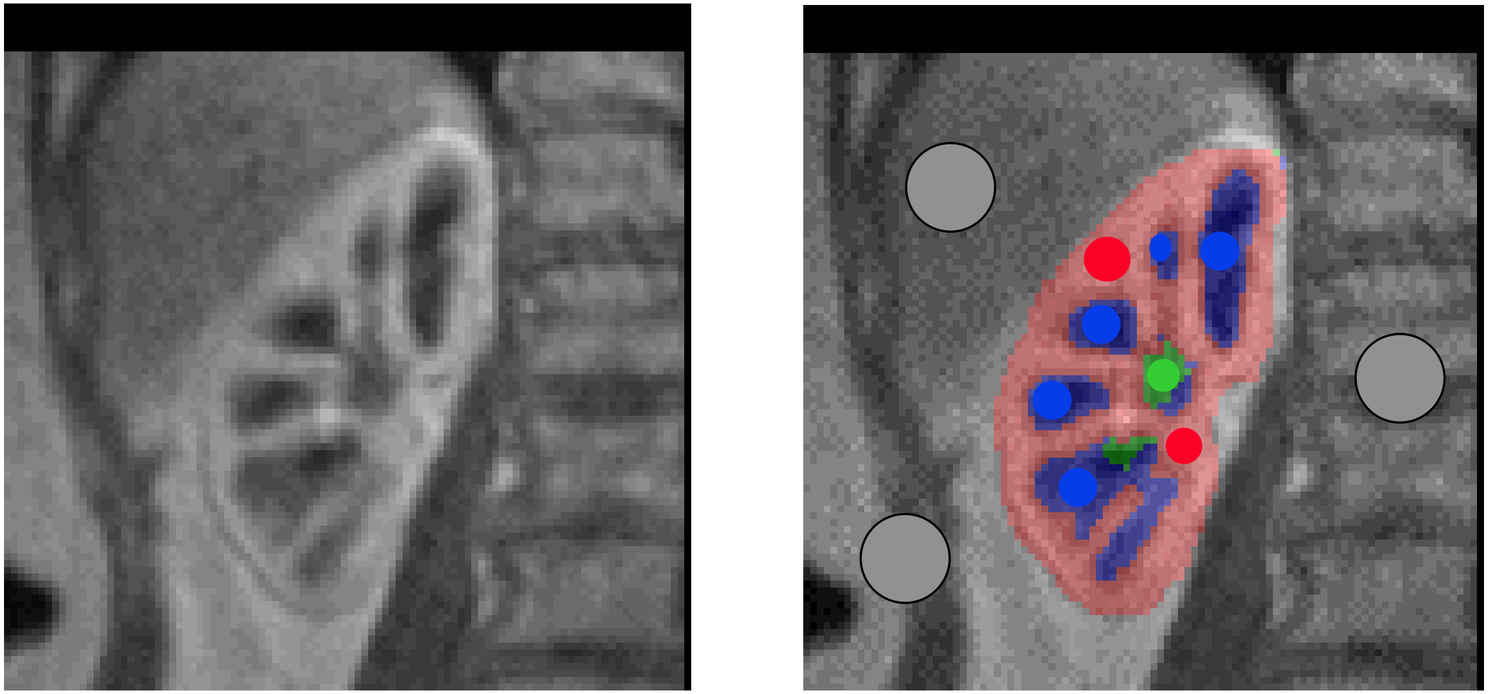

Obvious generalization of binary object extraction technique
(Boykov, Jolly, Funkalea 2004)

Multi-way graph cuts

Stereo/Motion with slanted surfaces

(Birchfield &Tomasi 1999)

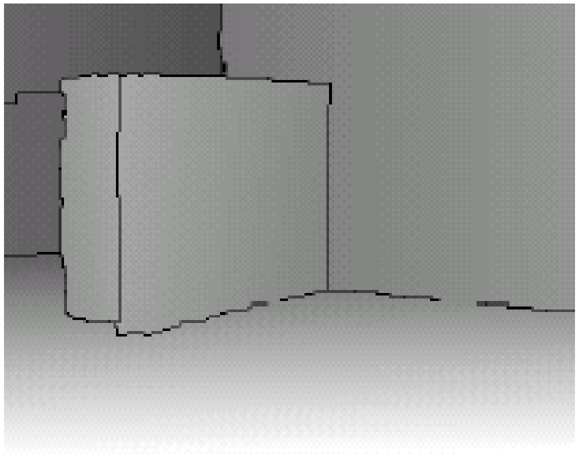

Labels = parameterized surfaces

EM based: E step = compute surface boundaries
M step = re-estimate surface parameters

Multi-way graph cuts

stereo vision

original pair of "stereo" images

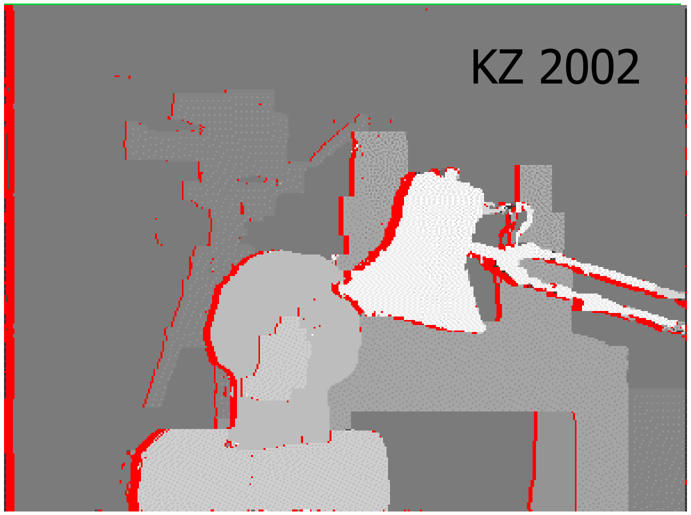

depth map

Multi-way graph cuts

Graph-cut textures

(Kwatra, Schodl, Essa, Bobick 2003)

similar to "**image-quilting**" (Efros & Freeman, 2001)

Multi-way graph cuts

Graph-cut textures

(Kwatra, Schodl, Essa, Bobick 2003)

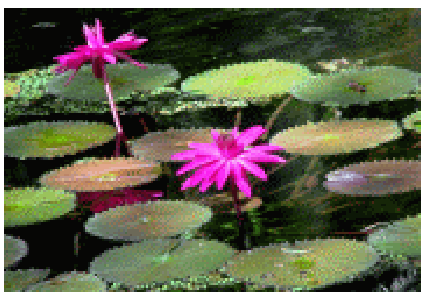

a-expansions vs. simulated annealing

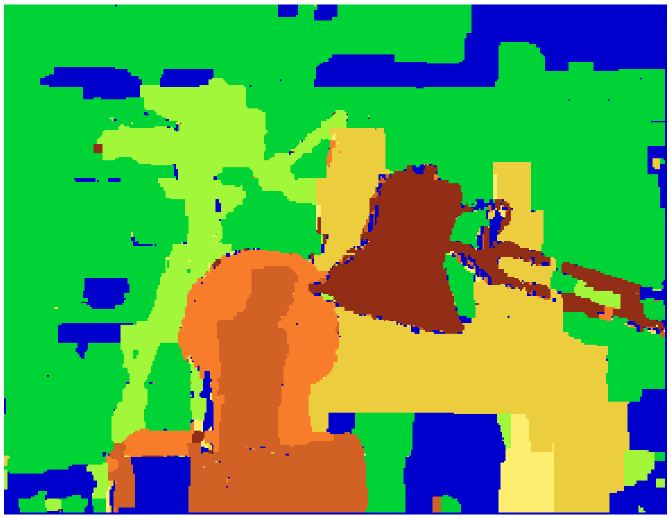

normalized correlation, start for annealing, 24.7% err

simulated annealing, 19 hours, 20.3% err

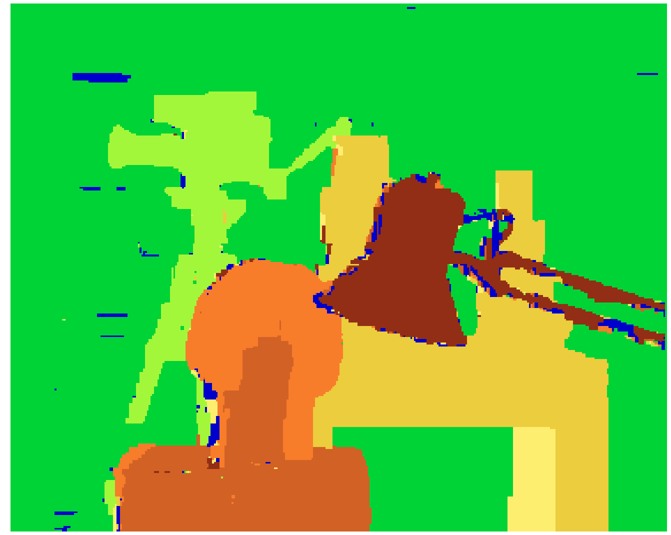

a-expansions (BVZ 89,01)
90 seconds, 5.8% err