edge detection in 3D

E. Meinhardt
V. Caselles & A. Frangi & G. Randall & E. Zacur

9–9–2006

Outline

Outline

Description of the problem

Goal of edge detection

To find the boundaries between the objects that appear in an image.

Edge detection *vs* Segmentation

- Related, but different, problems
- Different constraints (e.g. need of initialization)
- Different applications (matching versus picking)

Generalizing edge detection to 3D

Edges in 2D are curves

input

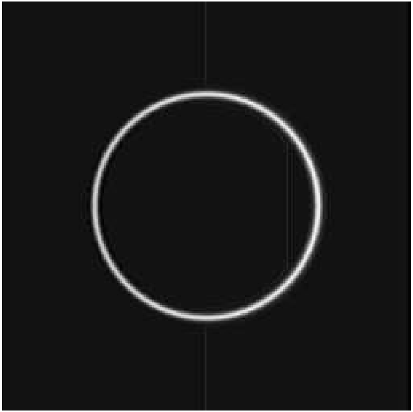

output

Generalizing edge detection to 3D

“Edges” in 3D are surfaces

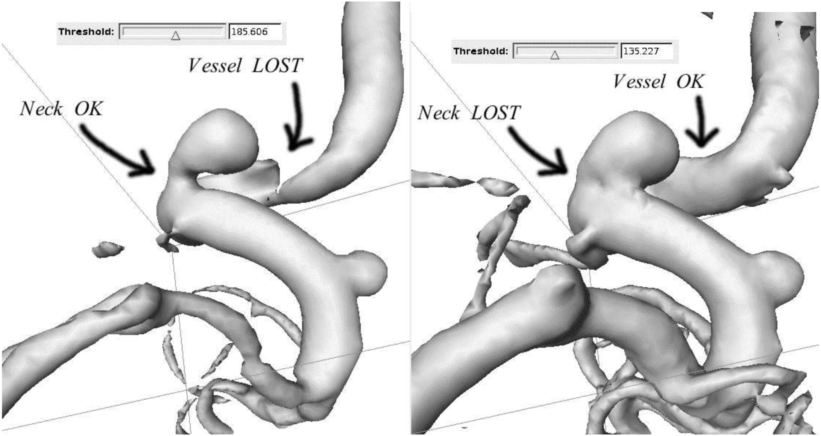

Edge detection in 2D

Point-wise methods

Differential operator that measures the “edgeness” of every point in the image.

- Easy to compute
- Trivially generalizable to 3D
- Set of disconnected points, difficult to use

Curve-based methods

Curve selection based on gestalt criteria by Desolneux et.al.

- Meaningful level lines
  - Easy to compute and generalize to 3D
  - Limited applicability
- Meaningful boundaries (pieces of level line)
  - Very good results
  - Not trivially generalizable to 3D

Our approach to 3D edge detection

General setting

- Define a family of subsets of the image that we are going to test for “edgeness”
- Define a test for edgeness of a subset

Our hypotheses

- The boundaries are formed by *large* pieces of level surfaces of the original grayscale image
- The boundaries have as high contrast as possible

Outline

Subsets that pass or not pass the test

Defining well-contrasted subsets

A contrario model for a subset

- We consider N subsets of independently selected points
- The distribution of their contrast is $H(\mu) = P(X_i \geq \mu)$

Test of good contrast

- Pick a set with n points and minimum contrast μ
- It passes the test whenever $NH(\mu)^n \leq 1$

Proposition

In the a-contrario model, the expected number of subsets that pass the test is ≤ 1.

A problem with the NFA defined using the minimum

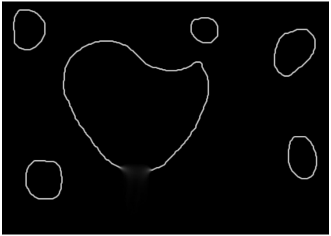

Alternative tests of good contrast

- We have defined a set $X_1, \ldots, X_n$ as meaningful when

$$\min\{X_1, \cdots, X_n\}$$

  is much larger than expected.
- But we could as well have used some statistic f other than the minimum:

$$f(X_1, \cdots, X_n)$$

  for example: the mean, the median, or the 10%th quantile.
- The computations are not harder.

A problem with the NFA of the 10% quantile

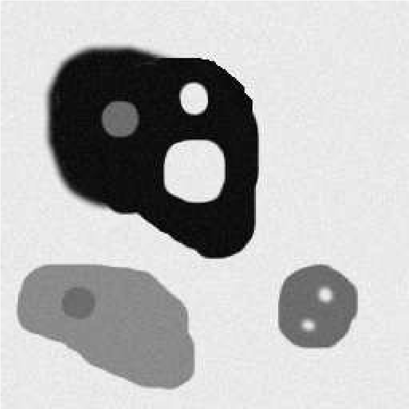

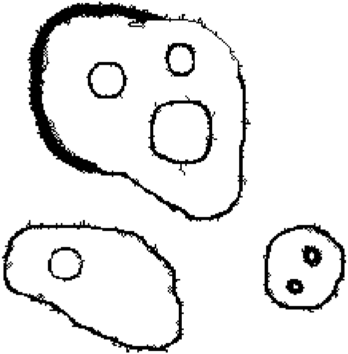

A problem with the NFA of the 10% quantile

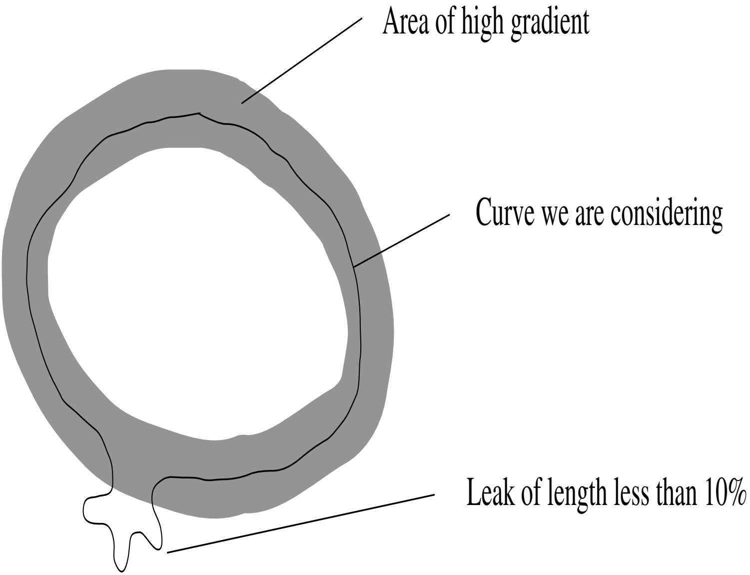

Outline

Meaningful level surfaces

Idea

Run the contrast test for all the level surfaces of the original image.

Implementation

- Easily implementable thanks to the *tree of shapes*
- Some numbers for a typical $100 \times 84 \times 72$ image:
  - 604800 voxels
  - 37768 level surfaces (nodes in the tree of shapes)
  - 2219 surfaces that pass the test
  - 272 surfaces having the best result in their branch

Discussion

- Direct generalization from 2D
- Best edges are not usually *whole* level surfaces

Some isosurfaces of a medical image

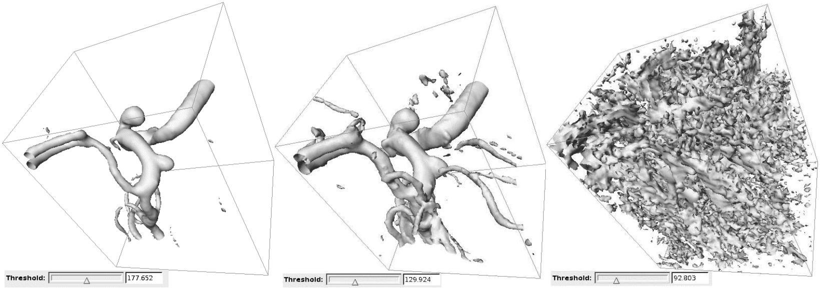

The contrast over a level surface

slice of a 3D image

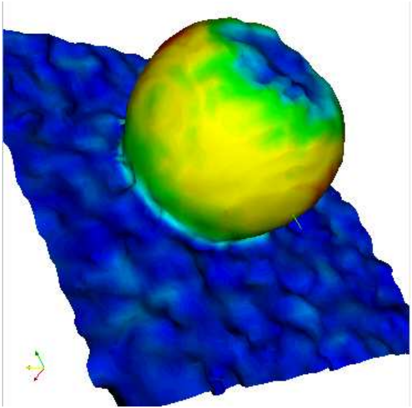

contrast over a level surface

General Mumford-Shah segmentation

The setting

- Ω an n-dimensional Riemanninan manifold
- V an m-dimensional Hilbert space
- $I : \Omega \to V$ a vector-valued function we want to approximate

The Mumford-Shah energy functional

Approximate I with a simplified version u that minimizes this energy:

$$E(u) = \int_{\Omega \setminus K} \|\nabla u\|_V \; + \; \lambda H^{n-1}(K) \; + \; \mu \int_{\Omega} \|I - u\|^2$$

where

- K is the discontinuity set of u, which is smooth on $\Omega \setminus K$
- λ, μ are scale parameters.

Piecewise constant Mumford-Shah segmentation

Energy of a partition

$$E(\text{Partition } \Omega_1, \ldots, \Omega_n) = \sum_i \int_{\Omega_i} \|I - m_i\|^2 + \lambda \sum_{i,j} l_{ij}$$

where

- m_i is the mean of I on Ω_i
- l_{ij} is the length of the border between Ω_i and Ω_j

Interpretation

- For $\lambda = 0$ we get the partition of the domain into the constant regions of I.
- For $\lambda = \infty$ we get the partition of the domain into one single piece.
- Increasing λ from 0 upwards we get a *hierarchy of partitions*, which can be organized in a tree.

The Mumford-Shah tree

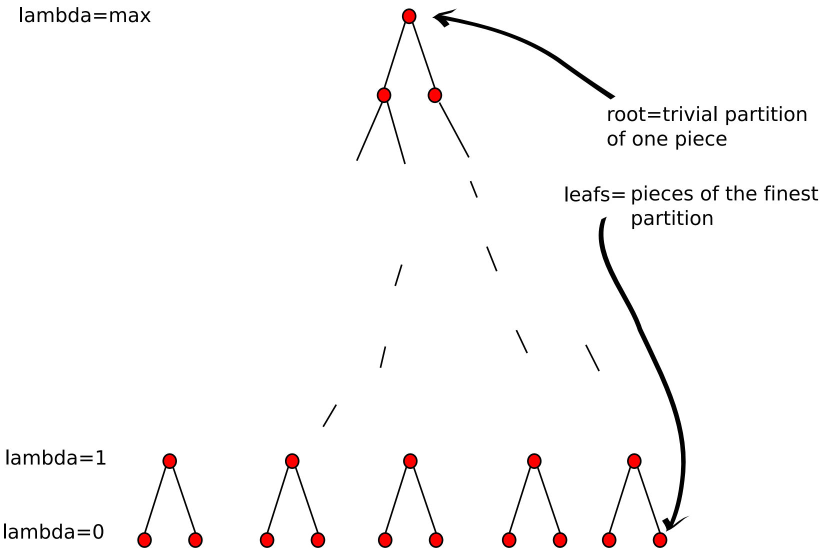

Example of Mumford-Shah segmentation

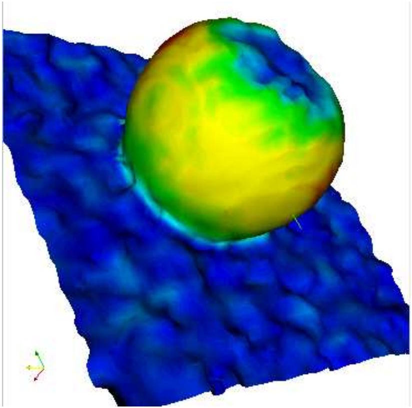

contrast over a level surface

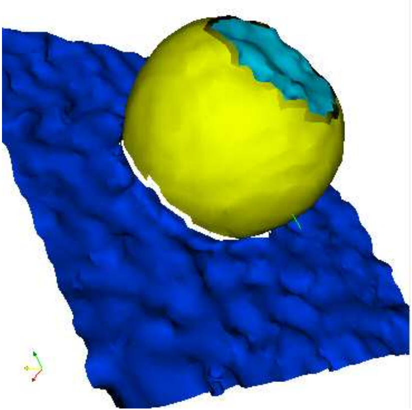

segmentation in 3 pieces

Use of the MS tree to define maximality

The hierarchy of partitions given by the Mumford-Shah tree allows us to define a disjoint set of pieces of each level surface that are maximally meaningful.

Conclusion

My two contributions

- Generalize the definition of the NFA for well-contrasted sets to statistics other than the minimum.
- Define maximal meaningful patches of 3D level surfaces in a usable way.

Current Work

- Finish the implementation of the MS tree
- Make the method useful for medical images:
  - Use other descriptors than the contrast to separate interfaces between organs and bones
  - Use an NFA with information learnt from manual segmentations to select only the “interesting” patches.