Image analysis with subpixel precision - The Coverage model

Coverage Segmentation

Nataša Sladoje and Joakim Lindblad
sladoje@uns.ac.rs, joakim@cb.uu.se

2012, Uppsala

The Coverage model
Nataša Sladoje and Joakim Lindblad
Coverage segmentation
Coverage segmentation by energy minimization

Outline

1. Coverage segmentation
   - Preliminaries
   - Soft thresholding
   - Graph based segmentation
   - Un-mixing based on local classification
2. Coverage segmentation by energy minimization
   - Energy terms
   - Minimization
   - Evaluation
   - Application example
   - Further work and extensions

The Coverage model
Nataša Sladoje and Joakim Lindblad
Coverage segmentation
Coverage segmentation by energy minimization

Coverage segmentation

Coverage segmentation methods address the task of extracting coverage information about objects in images, i.e., estimating a vector of coverage values for each image element.

Let $\mathbb{A}_m$ denote the set of m-component *segmentation vectors*

$$\mathbb{A}_m = \left\{ \boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_m) \in [0,1]^m \;\middle|\; \sum_{k=1}^{m} \alpha_k = 1 \right\}.$$

A *coverage segmentation* of an image I into m components is a set of ordered pairs

$$\left\{ \big((i,j), \boldsymbol{\alpha}(i,j)\big) \;\middle|\; (i,j) \in I_D, \boldsymbol{\alpha}(i,j) \in \mathbb{A}_m \right\}, \quad \alpha_k \approx \frac{|p_{(i,j)} \cap S_k|}{|p_{(i,j)}|},$$

where $S_k \subset \mathbb{R}^2$ is the extent of the k-th image component (object) and $I_D \subseteq \mathbb{Z}^2$ is the discrete image domain.

The continuous sets S_k are, in general, not known, and the values α_k have to be estimated from the image data.

The Coverage model
Nataša Sladoje and Joakim Lindblad
Coverage segmentation
Coverage segmentation by energy minimization

Coverage segmentation

Five methods which provide (approximate) coverage images:

1. Direct assignment of coverage values from a **continuous model**.
   - A. Tanács, C. Domokos, N. Sladoje, J. Lindblad, and Z. Kato. Recovering affine deformations of fuzzy shapes. SCIA 2009. LNCS-5575, pp. 735–744, 2009.
2. A method based on **mathematical morphology** and a double thresholding scheme.
   - N. Sladoje and J. Lindblad. High Precision Boundary Length Estimation by Utilizing Gray-Level Information. IEEE Trans. on PAMI, Vol. 31, No. 2, pp. 357–363, 2009.
3. A framework (and methods) for coverage segmentations of **graphs**.
   - F. Malmberg, J. Lindblad, I. Nyström. Sub-pixel segmentation with the image foresting transform. IWCIA 2009. LNCS-5852, pp. 201–211, 2009.
   - F. Malmberg, J. Lindblad, N. Sladoje, I. Nyström. A Graph-based Framework for Sub-pixel Image Segmentation. Theoretical Computer Science. Vol 412, No 15, pp. 1338-1349, 2011.
4. A method providing **local sub-pixel classification** extending any existing crisp segmentation.
   - N. Sladoje and J. Lindblad. Pixel coverage segmentation for improved feature estimation. ICIAP 2009. LNCS-5716, pp. 929-938, 2009.
5. An **energy based** method for regularized coverage segmentation.
   - J. Lindblad and N. Sladoje. Coverage Segmentation Based on Linear Unmixing and Minimization of Perimeter and Boundary Thickness. Pattern Recognition Letters. Vol 33, No. 6, pp. 728-738, 2012.

The Coverage model
Nataša Sladoje and Joakim Lindblad
Coverage segmentation
Coverage segmentation by energy minimization

Method 1: Use of a continuous segmentation model

From a continuous (crisp) representation it is, in general, straightforward to compute pixel coverage values, either analytically or numerically, e.g. based on supersampling.

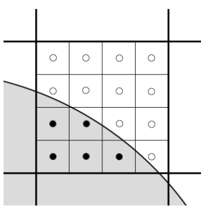

Approximate coverages: $\alpha_1 = \frac{5}{16}, \alpha_2 = \frac{11}{16}$

The Coverage model
Nataša Sladoje and Joakim Lindblad
Coverage segmentation
Coverage segmentation by energy minimization

Method 2: Soft thresholding and mathematical morphology

In many imaging situation, acquired pixel intensities correspond almost directly to pixel coverage values.

For example: Integration of photons over finite sized sensor elements, such as those of a digital camera.

- A reasonable model for low resolution images, where resolution is decided based on limited means for handling of the data, rather than the optical system. For example, this is often the case for low-resolution (surveillance) video.

However, noise may provide unreliable measurements. Appropriate pre-processing is recommended.

Soft thresholding and mathematical morphology

Properties of pixel coverage images

- The pixel coverage digitization leads to images where objects have grey edges which are never more than **one pixel thick** (if sampled at high enough resolution).

A pixel coverage segmentation method based on mathematical morphology in combination with a double thresholding scheme.

Given a grey-scale image, we seek a threshold couple, b and f, where pixels darker than b are defined to belong completely to the background, while pixels brighter than f belong completely to the foreground, such that the pixels in between form a not more than one pixel thick separating band.

In addition, we want the contrast between foreground and background, i.e., the difference $f - b$, to be as large as possible.

The requirement of a one pixel thin grey border is conveniently expressed using **grey-scale mathematical morphology**.

Algorithm

Input: A grey-scale image I.
Output: An approximates pixel coverage representation J with n positive grey-levels.

$b = 0; f = 0$
for each grey-level b'
  $F' = \{p \mid [\varepsilon I](p) > b'\}$ /* Foreground */
  if $F' \neq \emptyset$
    $f' = \min_{p \in F'} [\varepsilon \delta I](p)$
    if $f' - b' > f - b$ /* Better than previous */
      $f = f'; b = b'$
    endif
  endif
endfor

$n = f - b$

$$J(p) = \begin{cases} 0 & , \quad [\delta\varepsilon I](p) \leq b, \\ 1 & , \quad [\varepsilon\delta I](p) \geq f, \\ \frac{I(p)-b}{n} & , \quad \text{otherwise.} \end{cases}$$

Illustrative example – Matlab "coins.png"


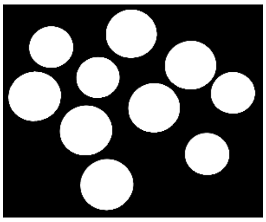


Illustrative example – Matlab "coins.png"

To do: Local thresholding


Local thresholding (right) preserves coverage information better.

Method 3: Sub-pixel Segmentation with the Image Foresting Transform

Work with Filip Malmberg

- The Image Foresting Transform (IFT) is a framework for supervised and often interactive segmentation.
- Given a set of seed-points with user-defined labels, the IFT completes the labelling by computing minimal cost paths from all image elements to the seed-points.
- Each image element is given the (crisp) label of its closest seed-point.
- We propose a modified version of the IFT that allows mixed labels at region boundaries.

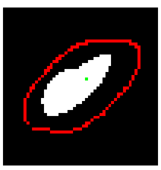
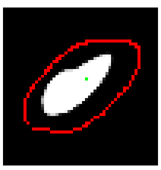

Sub-pixel Segmentation with the Image Foresting Transform

1. Compute the standard (crisp) IFT.
2. For each edge connecting two nodes (voxels) with different label, compute the intersection point.
3. Assign mixed labels to the nodes by integrating the labels of its associated edges.

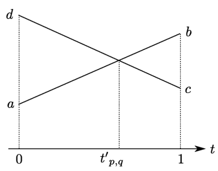

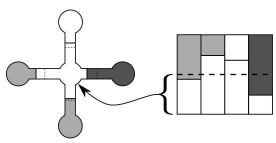

Sub-pixel Segmentation with the Image Foresting Transform

The proposed sub-pixel segmentation provides **much more stable feature estimates** than the original IFT.

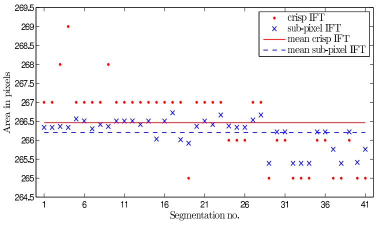

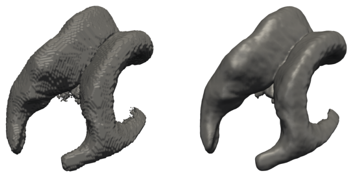

Lateral ventricles

Generalization

- Coverage values estimated from any fuzzy segmented graph

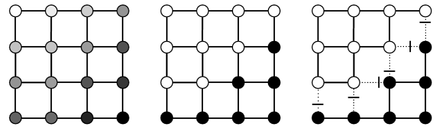

Figure 3: Computing located cuts in the context of defuzzification. (Left) A fuzzy vertex segmentation $\mathcal{V}$ of a graph. (Middle) Corresponding defuzzified vertex segmentation $\hat{\mathcal{V}}$. (Right) Information from $\mathcal{V}$ and $\hat{\mathcal{V}}$ is combined to find a location $\mathcal{T}$ such that $(\partial\hat{\mathcal{V}}, \mathcal{T})$ is a located cut.

Method 4: Un-mixing based on local classification

Assumption
Partial pixel coverage exist only at the object boundaries of an existing crisp segmentation.

Approach
Re-assign class belongingness to the boundary pixels based on a local classification using the surrounding non-boundary pixels.

To obtain a pixel coverage segmentation, we propose a method composed of the following four steps:

1. Application of a crisp segmentation method, appropriately chosen for the particular task
2. Selection of pixels to be assigned partial coverage
3. Application of a liner mixture model for "de-mixing" of partially covered pixels and assignment of pixel coverage values
4. Ordered thinning of the set of partly covered pixel to provide one pixel thin 4-connected regions of mixed pixels

Steps 1 and 2.

1. **Any** crisp segmentation model.
   - For the example to come, we used Linear Discriminant Analysis in combination with Iterated Relative Fuzzy Connectedness[1]
2. **Selection of pixels to re-evaluate**
   - All pixel which are 4-connected to a pixel with a different label.

[1] J. Lindblad, N. Sladoje, V. Ćurić, H. Sarve, C.B. Johansson, and G. Borgefors. Improved quantification of bone remodelling by utilizing fuzzy based segmentation. SCIA 2009

Step 3. Computation of partial pixel coverage values

3.1 Estimate the spectral properties c_k of the pure classes locally.

- The mean values of the respective classes present in the assumed **completely covered** pixels in a local Gaussian neighbourhood.

3.2 Compute the mixture proportions a_k of the pixels selected in step 2.

- The intensity values of a mixed pixel $p = (p_1, p_2, \ldots, p_n)$ (n being the number of channels of the image) are assumed, in a noise-free environment, to be a convex combination of the pure classes c_k:

$$p = \sum_{k=1}^{m} \alpha_k c_k\,, \quad \sum_{k=1}^{m} \alpha_k = 1\,, \quad \alpha_k \geq 0\,. \tag{1}$$

where each coefficient α_k corresponds to the coverage of the pixel p by an object of a class c_k.

Step 3. Computation of partial pixel coverage values

In the presence of noise, it is not certain that there exists a (convex) solution to the linear system (1). Therefore we reformulate the problem as follows:

Find a point p^* of the form $p^* = \sum_{k=1}^{m} \alpha_k^* c_k$, such that p^* is a *convex* combination of c_k and the distance $d(p, p^*)$ is minimal. We solve the constrained optimization problem by using Lagrange multipliers, and minimize the function

$$F(\alpha_1, \ldots, \alpha_m, \lambda) = \left\| p - \sum_{k=1}^{m} \alpha_k c_k \right\|_2^2 + \lambda\left(\sum_{k=1}^{m} \alpha_k - 1\right)$$

over all $\alpha_k \geq 0$. This leads to a least squares type of computation.
The obtained solution provides estimated partial coverage of the pixel p by each of the observed classes c_k.

Step 4. Ordered thinning

To ensure one pixel thick boundaries, the “least” mixed pixels are one at a time assigned to their most prominent class, until only one pixel thick mixed boundaries remain.

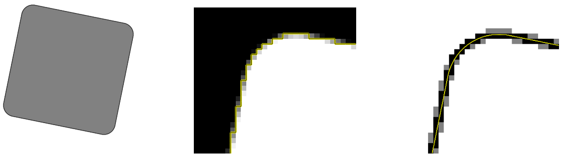

(a) Test object (b) Part of pixel coverage segm. (c) Part of re-evaluated set

Quantitative evaluation

- noise sensitivity

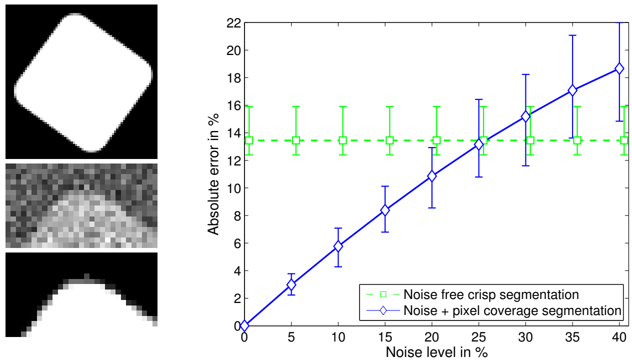

Left: (top) Synthetic test objects. (middle) Part of object with 30% noise added. (bottom) Coverage segmentation result for 30% noise. **Right:** Average absolute error of coverage values of object *border pixels* for different noise levels. Lines show averages for 50 observations and bars indicate max and min errors.

Illustrative example – Matlab ”coins.png”

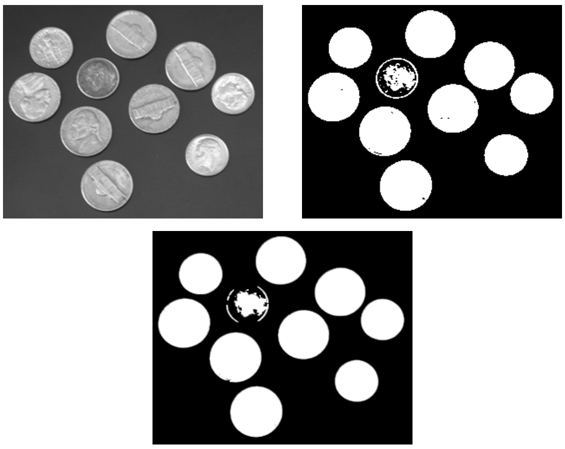

Otsu thresholding fails; garbage in, garbage out.

Illustrative example – Matlab ”coins.png”

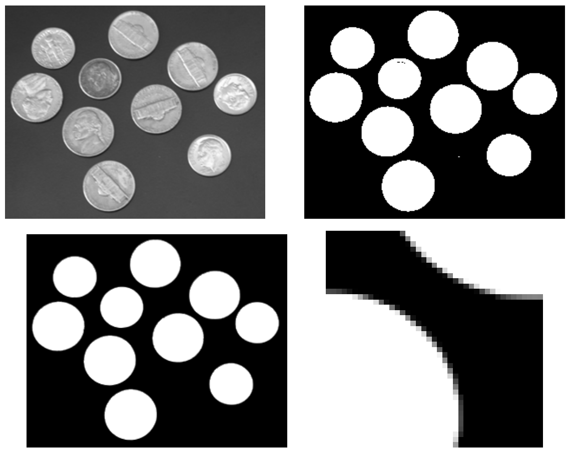

A better threshold.