From Contours to Regions: An Empirical Evaluation

Pablo Arbeláez[1], Michael Maire[1], Charless Fowlkes[2], and Jitendra Malik[1]

[1]*University of California at Berkeley - Berkeley, CA 94720* [2]*University of California at Irvine - Irvine, CA 92697*

Abstract

We propose a generic grouping algorithm that constructs a hierarchy of regions from the output of any contour detector. Our method consists of two steps, an Oriented Watershed Transform (OWT) to form initial regions from contours, followed by construction of an Ultrametric Contour Map (UCM) defining a hierarchical segmentation. Extensive experimental evaluation demonstrates that, when coupled to a high-performance contour detector, the OWT-UCM algorithm produces state-of-the-art image segmentations. We achieve an F-measure of 0.71 on the Berkeley Segmentation Dataset Benchmark, a significant improvement over Mean Shift (0.63) and Normalized Cuts (0.62).

Hierarchical Segmentation from Contours

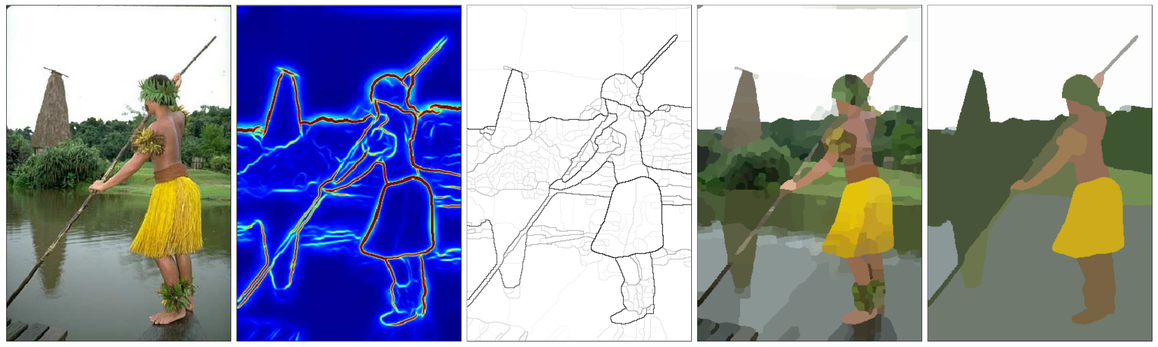

Contour Detector **OWT-UCM** **Hierarchical Segmentation**

Oriented Watershed Transform (OWT)

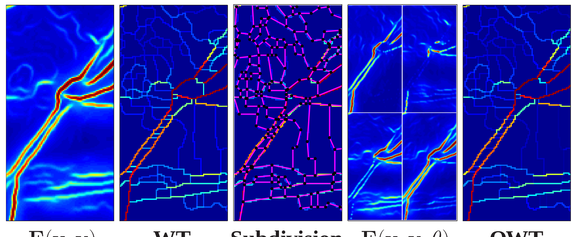

$\mathbf{E(x,y)}$ **WT** **Subdivision** $\mathbf{E(x,y,\theta)}$ **OWT**

Form initial regions (oversegmentation) using contour detector:

1. Consider detector whose output $E(x, y, \theta)$ indicates strength of boundary at pixel (x, y) and orientation θ.
2. Compute $E(x, y) = \max_\theta E(x, y, \theta)$.
3. Apply watershed transform (WT) producing:
   - Catchment basins $\mathcal{P}_0$ as regions
   - Arcs $\mathcal{K}_0$ as boundaries
4. Subdivide arcs $\mathcal{K}_0$ into approximating straight line segments.
5. Obtain OWT by reweighting arcs according to $E(x, y, \theta)$ at the orientation θ of their associated line segments.

Ultrametric Contour Map (UCM)

The UCM defines a duality between closed, non-self-intersecting weighted contours and a hierarchy of regions. We construct this hierarchy using a graph-based region merging algorithm:

1. Construct graph $G = (\mathcal{P}_0, \mathcal{K}_0, W(\mathcal{K}_0))$ given by OWT
2. Iteratively merge regions by removing min weight boundary.
3. Produces region tree where:
   - Root is entire image
   - Leaves are $\mathcal{P}_0$
   - $Height(R)$ is boundary threshold at which R first appears
   - $Distance(R_1, R_2) = min\{Height(R) : R_1, R_2 \subseteq R\}$

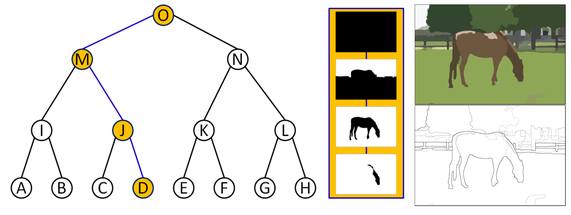

Software is available on our website: http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/

Boundary Benchmarks

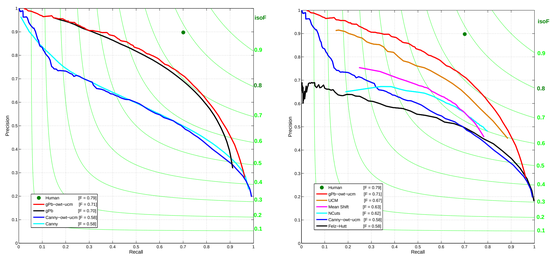

Berkeley Segmentation Dataset (BSDS) Benchmark. Left: The OWT-UCM algorithm preserves contour detector quality. **Right:** Paired with gPb [Maire et. al. 08] as input, gPb-owt-ucm outperforms existing segmentation approaches.

Results

gPb-owt-ucm **ODS** **OIS**

Region Benchmarks

Segment Covering

Overlap between two regions R and R':

$$\mathcal{O}(R, R') = \frac{|R \cap R'|}{|R \cup R'|}$$

Covering of groundtruth segmentation S by a segmentation S':

$$\mathcal{C}(S' \to S) = \frac{1}{N} \sum_{R \in S} |R| \cdot \max_{R' \in S'} \mathcal{O}(R, R')$$

Probabilistic Rand Index [Unnikrishnan et. al. 07] [Yang et. al. 08]

Given a set of ground-truth segmentations $\{G_k\}$,

$$PRI(S, \{G_k\}) = \frac{1}{T} \sum_{i<j} [c_{ij} p_{ij} + (1 - c_{ij})(1 - p_{ij})]$$

where c_{ij} is the event that pixels i and j have the same label and p_{ij} its probability.

Variation of Information [Meila 05]

Distance between two clusterings of data C and C' given by

$$VI(C, C') = H(C) + H(C') - 2I(C, C')$$

Here C and C' are test and ground-truth segmentations.

Segmentation Method	ODS	OIS	Best	PRI	VI
human	0.73	0.73	–	0.87	1.16
gPb-owt-ucm	**0.58**	**0.64**	**0.74**	**0.81**	**1.68**
Mean Shift	0.54	0.58	0.64	0.78	1.83
Felz-Hutt	0.51	0.58	0.68	0.77	2.15
Canny-owt-ucm	0.48	0.56	0.67	0.77	2.11
NCuts	0.44	0.53	0.66	0.75	2.18

Region Benchmarks on the BSDS. For each algorithm, we report the score of the covering of ground-truth segments according to optimal dataset scale (ODS), optimal image scale (OIS), or Best covering criteria, as well as the PRI and VI.

Interactive Segmentation

Segmentation trees generated by the OWT-UCM algorithm provide an alternative to the graph-cuts formalism for user-assisted refinement:

1. User paints labels on a subset of regions.
2. Partial labeling is extended to a full one by assigning to each unlabeled region the label of its closest (in terms of the ultrametric distance) labeled region.
3. Additional labels placed interactively until reaching desired result.

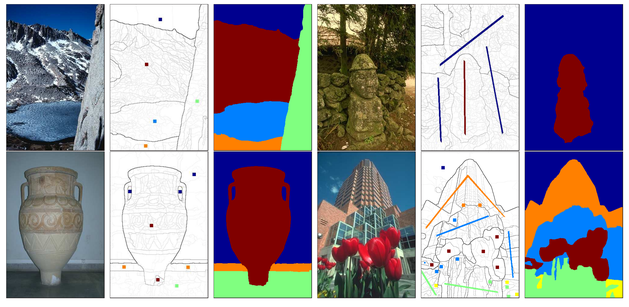

This work was partially supported by ONR MURI N00014-06-1-0734.