

Weakly-Supervised Segmentation of Non-Gaussian Images via Histogram Adaptation

Jonas August (jonas@cs.cmu.edu) and Takeo Kanade (tk@cs.cmu.edu)

Robotics Institute, Carnegie Mellon University, Pittsburgh, PA

Abstract. Here we apply an active contour model that allows for arbitrary intensity distributions inside and outside the boundary of an object to be segmented in an image. Computationally, we estimate intensity histograms both inside and outside the current boundary estimate, and use these histograms to define an image energy as their log-likelihood ratio. Training the model with accurate example segmentations is unnecessary; initialization with a crude, user-provided segmentation is sufficient.

Speckle and other forms of noise common in medical images is non-Gaussian, and yet active contour and level-set techniques applied to segment these images have assumed Gaussian noise [1]. To overcome this limitation, we allow for arbitrary intensity distributions inside and outside the boundary of an object to be segmented.

A key barrier to applying non-Gaussian statistical models is parameter estimation. We take a weakly-supervised approach in which region statistics from a crude user-provided initial segmentation are adaptively updated as a level-set method updates the contour. Specifically, we obtain the histograms of pixel values both inside and outside of the user-provided segmentation. These histograms define an external or image energy which is combined with an internal energy that penalizes length in an active contour method. Gradient descent on the combined energy gives rise to an updated segmentation, leading to updated interior and exterior intensity histograms, which in turn give rise to an updated combined energy functional. This process is iterated until convergence. We now provide mathematical background.

In a probabilistic expression of the active contour method [2], one seeks that curve \mathcal{C} minimizing the combined external and internal energy

$$E(\mathcal{C}) = - \int_{\mathcal{R}} \log p(I(\mathbf{x})|\mathbf{x} \in \mathcal{R})d\mathbf{x} - \int_{\Omega \setminus \mathcal{R}} \log p(I(\mathbf{x})|\mathbf{x} \notin \mathcal{R})d\mathbf{x} + \beta \text{length}(\mathcal{C}),$$

where $\Omega \subset \mathbb{R}^2$ is the entire image region; \mathbf{x} is a point in \mathbb{R}^2 ; $I = I(\mathbf{x})$ is one of a finite number of intensities; $\mathcal{R} \subset \Omega$ is a segmented region with boundary $\partial\mathcal{R} = \mathcal{C}$; $\log p(I(\mathbf{x})|\mathbf{x} \in \mathcal{R})$, the interior likelihood, (resp., $\log p(I(\mathbf{x})|\mathbf{x} \notin \mathcal{R})$, the exterior likelihood) is the natural logarithm of the probability that the intensity at \mathbf{x} is $I(\mathbf{x})$, given that \mathbf{x} is inside (resp., outside) \mathcal{R} ; and β is a positive constant. Let $h_{\text{in}}(I)$ and $h_{\text{out}}(I)$ denote the (normalized) intensity histograms inside and

outside \mathcal{C} . We use $h_{\text{in}}(I)$ (resp., $h_{\text{out}}(I)$) as an estimate of the *unknown* interior (resp., exterior) likelihood, because if \mathcal{C} is the true boundary, then $h_{\text{in}}(I)$ (resp., $h_{\text{out}}(I)$) is a maximum likelihood estimate (subject to a normalization constraint) of $p(I|x \in \mathcal{R})$ (resp., $p(I|x \notin \mathcal{R})$). The curve evolution we use as an approximate gradient descent of $E(\mathcal{C})$ is:

$$\frac{\partial \mathcal{C}}{\partial t} = \left(\log \frac{h_{\text{in}}(I(\mathbf{x}))}{h_{\text{out}}(I(\mathbf{x}))} - \beta \kappa \right) \mathcal{N},$$

where t is time in the evolution, κ is curvature, and \mathcal{N} is the outward normal. The second term $\beta \kappa$ is the familiar length-shortening term from the geometric heat equation. The first term, a log-likelihood ratio, is a *non-homogeneous* area term, which expands the curve where the interior is more likely ($h_{\text{in}} > h_{\text{out}}$), and shrinks it where the exterior is more likely. The log-likelihood ratio replaces the artificial balloon forces common in edge-based active contours. The level-set method implemented in the Insight Toolkit was used for computations (Fig. 1).

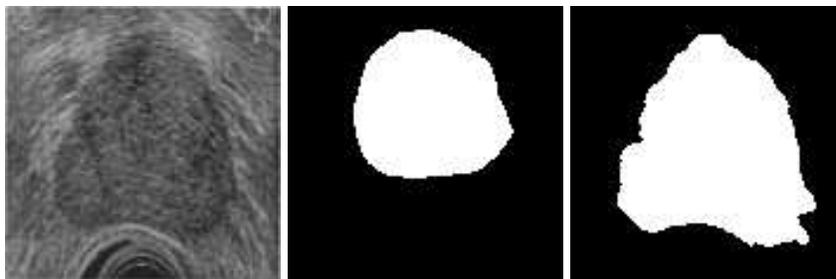


Fig. 1. Prostate segmentation. (Left) Ultrasound image of prostate; (center) crude initial segmentation from which initial interior and exterior statistics were determined by histogramming corresponding portions of image; (right) final segmentation after convergence (1800 iterations of level set method). The bottom right protrusions are consistent with the image, suggesting a stronger shape prior may be required.

Although we have demonstrated the log-likelihood area term using intensity histograms, we can obtain log-likelihood ratios for other image properties such as gradient magnitude, Gabor filter responses, etc., indeed any feature that one would use to distinguish the interior texture properties from those of the exterior. This and the extension to 3-d images shall be described in our future papers.

References

1. T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. Image Proc.*, 10(2):266–277, 2001.
2. S. C. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and bayes/mdl for multiband image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18:884–900, 1996.