

Nonparametric Intensity Priors for Level Set Segmentation of Low Contrast Structures

Sokratis Makrogiannis¹, Rahul Bhotika¹, James V. Miller¹, John Skinner Jr.²,
and Melissa Vass²

¹ GE Global Research, One Research Circle, Niskayuna, NY 12309, USA
{makrogi, bhotika, millerjv}@research.ge.com

² GE Healthcare, 3000 N. Grandview Blvd., Waukesha, WI 53188, USA
{melissa.vass, john.skinner}@ge.com

Abstract. Segmentation of low contrast objects is an important task in clinical applications like lesion analysis and vascular wall remodeling analysis. Several solutions to low contrast segmentation that exploit high-level information have been previously proposed, such as shape priors and generative models. In this work, we incorporate *a priori* distributions of intensity and low-level image information into a nonparametric dissimilarity measure that defines a local indicator function for the likelihood of belonging to a foreground object. We then integrate the indicator function into a level set formulation for segmenting low contrast structures. We apply the technique to the clinical problem of positive remodeling of the vessel wall in cardiac CT angiography images. We present results on a dataset of twenty five patient scans, showing improvement over conventional gradient-based level sets.

1 Introduction

This paper introduces a method for segmenting low contrast regions in Computed Tomography (CT) volumes by integrating local nonparametric intensity statistics into the level set framework. The drive for early detection and quantification of disease has greatly improved the spatial resolution and sensitivity of CT scanners. This has led to the growing need for methods for sub-pixel accurate segmentation and boundary delineation of small structures with volumes ranging from 25 mm^3 to 500 mm^3 , or equivalently, 200 to 4000 voxels. Segmentation of such small structures is further complicated by low contrast, partial volume averaging, and other imaging artifacts that make it difficult to robustly detect the boundary between the object of interest and its background and surrounding structures (see Fig. 1). Furthermore, the underlying anatomical structures of interest often have intensity distributions that are not robustly captured by standard parametrizations or lower-order sufficient statistics, as shown in Fig. 2. These factors motivate the importance of continued investigation into methods for improving the segmentation of small, low-contrast structures with spatially varying intensity information.




Fig. 1. Examples of low contrast structures in CT images. The image on the left shows a coronary artery with a thickened wall barely visible against the background cardiac tissue. The image on the right shows a subsolid nodule that appears “translucent” and diffuses into surrounding lung tissue.




Fig. 2. Left: Example of ramp intensity profile (yellow) and corresponding gradient magnitude (blue) that causes gradient-based level set to propagate beyond the real boundary given by the reference segmentation (green). Right: Histograms of pixel intensities in vessel wall remodeling regions (blue) and in nearby cardiac tissue and background (red) were computed over 20 subjects. The histograms overlap considerably, but have overall distinctive distributions. Note also that the cardiac tissue distribution is not a Gaussian and difficult to parametrize.

Previously, approaches using parametric deformable surface models that propagate by minimizing an energy functional which includes internal (stretching, bending), image (usually gradient-based), and external (user-defined or derived from priors) forces have been proposed [1,2]. These often produce accurate segmentation but have limitations with respect to topological flexibility and the segmentation of complex structures. These limitations are somewhat overcome by level set methods, which represent the evolving surface implicitly. Early techniques such as fast marching supported only monotonic motion [3] and were followed by more sophisticated approaches such as shape-based segmentation [4], geodesic active contours [5], or methods using prior shape statistics to regularize the level set formulation [6,7]. A solution for the piecewise constant case of the Mumford-Shah functional [8] in the level set context was presented in [9],

but only accounted for first-order foreground and background statistics, approximated using average intensities. In [10], a data consistency term was used to incorporate local intensity information, but the method makes several assumptions: the overall image intensity distribution is assumed to be bimodal such that it can be factored into underlying foreground and background distributions which themselves are assumed to be unimodal, and whose parameters can be recovered from an initial coarse segmentation. Recently, the Bhattacharyya distance, a nonparametric dissimilarity measure, was integrated into the level set framework [11], but the method does not allow the use of prior statistics to drive the segmentation process.

This paper introduces a technique that makes use of nonparametric prior models of intensity within the level set variational framework to accurately segment low contrast regions. Reference segmentations provided by clinical specialists were used to build nonparametric models of the intensity distribution of the object of interest and its surrounding background. During segmentation, we first estimate the dissimilarity of the local intensity distribution in the vicinity of each voxel in the image to the two model distributions using Earth Mover's Distance (EMD) or the Mallows Distance [12,13] as a local indicator function for the likelihood of the voxel belonging to the foreground. The local indicator function is integrated with gradient-based information to define a new speed term in the level set partial differential equation. There are two major contributions of this work. First, it incorporates and balances both boundary and texture information into the segmentation process, while retaining the topological flexibility, sub-pixel accuracy, and boundary extraction capabilities of the level set framework. Second, it uses prior statistical knowledge about the object and background intensity distributions without imposing a specific parametrization, e.g., a Gaussian distribution, on the class models. In the remainder of this paper, we provide the theoretical background, formulate our approach, present a clinical application, and show experimental results.

2 Background: The Level Set Formulation

In level set approaches the propagating surface $S(t)$ is represented by the zero level set of a higher dimensional function $\psi(\omega(t); t) : \mathbb{R}^3 \rightarrow \mathbb{R}$ that evolves with time in the Eulerian coordinate system, and may be defined as $S(t) \equiv \psi(\omega(t); t) = 0$. The main benefits of the level set formulation are topological flexibility and differentiability [3,4]. The basic equation of level set propagation is given by:

$$\frac{\partial \psi}{\partial t} + F(\omega(t); t) \cdot |\nabla \psi| = 0, \quad (1)$$

where $F(\cdot)$ is the speed function that controls the evolution of the level set and $\omega(t)$ denotes the spatial coordinates in the d -dimensional space $\omega \in N^d$ that varies with time t . Given an initial position of the surface, $S(t = 0)$, a signed distance transform is typically applied to generate $\psi(\omega(t = 0); t = 0)$. After setting the initial conditions, the evolution of $\psi(\omega(t); t)$ and deformation of $S(t)$ are obtained by solving (1) at each time point t .

In the conventional level set framework, surface evolution is governed by geometric and image-based features, such as the surface curvature κ , the gradient magnitude of the image intensity I , and an external (inflation/deflation) driving force c [4]. This can be expressed as

$$F_G(\omega(t); t) = [c + \kappa(\omega(t); t)] \cdot g(\nabla I) \cdot \zeta(P), \quad (2)$$

where $g(\nabla I) = 1 / (1 + |\nabla I * G_\sigma|)$ with G_σ being a Gaussian with standard deviation σ used to smooth out noise. P is an optional spatial prior term that may be used if prior shape, size, or location information is available. Also, $\zeta(X)$ is the sigmoid function that normalizes the dynamic range of a feature X and is given as $\zeta(X) = X_{\min} + (X_{\max} - X_{\min}) / \left(1 + \exp\left(-\frac{X - \beta}{\alpha}\right)\right)$. When the boundary between the object and background is diffuse, as is often the case in biomedical imaging, the edge profile follows a ramp model (see Fig. 2 (left)). In these cases, gradient-based implementations as in (2) tend to propagate over the boundary, leaking into the background. More accurate segmentation can be achieved by using prior statistical information about the intensity distribution of an object and its surrounding area to complement the gradient information. However, first and second order statistics or parametric representations of distributions are often not sufficient for separation of the object from the background, as shown in Fig. 2 (right). We propose the integration of higher order intensity statistics into the level set formulation via a nonparametric dissimilarity measure for segmenting low contrast structures.

3 Incorporating Nonparametric Priors in Level Sets

A local indicator function, defined at each voxel and denoted by $H(\omega)$, that expresses the likelihood of the voxel belonging to a foreground object can be integrated in the level set formulation along with gradient magnitude and curvature to define a new speed term F_H as follows (with the help of (2)):

$$F_H(\omega(t); t) = F_G(\omega(t); t) \cdot \zeta(H), \quad (3)$$

where $\zeta(\cdot)$, again, is the sigmoid function. So far, the only requirement imposed on $H(\cdot)$ is that it can be evaluated locally. Next, we define and derive a form of $H(\cdot)$ in terms of a dissimilarity measure of a locally sampled intensity distribution to nonparametric intensity distributions of foreground and background classes established from exemplars.

We represent the prior statistical information of foreground and background regions using normalized histograms of intensity. The histogram of foreground regions was learned from clinical image datasets where the foreground regions were delineated by clinical imaging experts. The histogram of the background was derived from the same datasets by automatically selecting regions surrounding the delineated foreground voxels. Fig. 2 (right) shows the histograms for remodeled sections of coronary vessel wall and the nearby cardiac tissue over a set of twenty CT volumes.

The histograms of the foreground and background intensity models are used to construct the indicator function $H(\cdot)$ in (3). At each voxel ω , a histogram built from a small neighborhood about ω is compared to the model foreground and background histograms using the Earth Mover's Distance [12,13] as a well defined norm on the differences between two distributions

$$D^2(\omega; M_i, \mathcal{N}(\omega)) = \sum_k [\text{CDF}(k, M_i) - \text{CDF}(k, \mathcal{N}(\omega))]^2, \quad (4)$$

where M_i is a model (foreground or background), $\mathcal{N}(\omega)$ is a neighborhood about ω , and $\text{CDF}(\cdot)$ is the cumulative probability distribution of intensity for a model or over a neighborhood. The dissimilarity measures of the sample histogram over $\mathcal{N}(\omega)$ to each model histogram are combined into an indicator function given by

$$H(\omega) = D^2(\omega; M_1, \mathcal{N}(\omega)) - D^2(\omega; M_2, \mathcal{N}(\omega)) - T \quad (5)$$

where M_1 and M_2 denote the background and object model respectively, and T represents a difference between the prior probabilities of the two models. $H(\cdot)$ can be plugged into (3) to give the nonparametric intensity-distribution weighted speed function, F_H .

4 Clinical Application: Vessel Wall Segmentation

The proposed algorithm was tested on cardiac CT angiography (CTA) datasets in the context of vessel wall delineation for measuring positive remodeling. CTA is a minimally invasive imaging procedure used to examine the health of blood vessels in key areas of the human body such as brain, heart, lung, abdomen, and kidneys. It involves injection of a contrast agent, usually iodine-based, to enhance the image contrast inside the vessel lumen. This enables the segmentation of the lumen using simple algorithms such as morphological operations or region growing and is useful for detecting stenoses, i.e., narrowing of the vessel lumen.

Positive wall remodeling, on the other hand, is defined as an increase in the vessel wall thickness, resulting in the bulging out of the outer boundary of the vessel wall. This is usually observed at atherosclerotic lesion sites and often indicates an unstable clinical presentation associated with high risk of an acute cardiac event [14,15]. Accurate segmentation plays a significant role here because the coronary vessel walls are thin, typically 2-5 pixels wide in state-of-the-art CT. Unfortunately, compared to the vessel lumen, the image contrast of the vessel wall is quite low. This is due to surrounding cardiac tissue, partial voluming, and noise. Thus, accurate localization of the wall boundary using gradient-based level sets is not possible. On the other hand, the intensity distribution of the positive remodeling voxels marked by clinical specialists was observed to differ from that of its surrounding tissue and background, as shown earlier in Fig. 2. Therefore, the proposed speed function given in (3) which incorporates nonparametric intensity statistics should yield improved segmentation.

5 Experimental Results

Preprocessing steps that are not the focus of this paper were used to isolate the bright contrast-filled vessel lumen regions in CTA datasets. For evaluating the proposed algorithm, we operated only on the vessel segments containing examples of wall remodeling. A volume equal to three times the nominal diameter of the vessel lumen was fixed as the domain for the level set propagation. A signed distance function was computed from the lumen surface and the level set was initialized at 0.5 mm outside the lumen.




Fig. 3. Segmentation results. Row (a) Original CT intensity data, (b) reference standard showing positive remodeling regions overlaid on the CT data, (c) segmentation using conventional gradient-based level sets (G-LS), and (d) segmentation results using the proposed histogram-based level set (H-LS) approach. The first two columns show axial and coronal views. For better visualization, the last two columns show reformatted views across and along the vessel. G-LS propagates over the wall boundary into the background areas, while H-LS produces more accurate delineation.

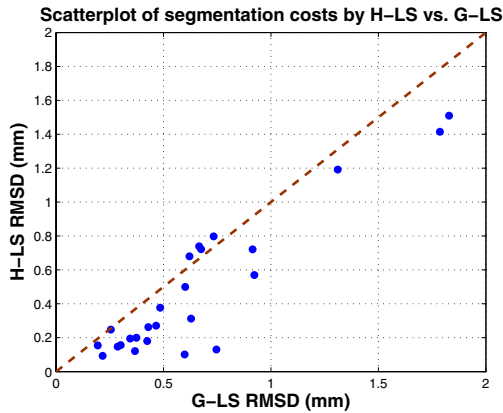


Fig. 4. Scatterplot of segmentation costs produced by H-LS versus G-LS for 25 patient CT scans (in blue). The segmentation cost is the root mean squared distance (RMSD) of the segmented surface from the reference surface. Points below the diagonal red line correspond to lower RMSD values of H-LS compared to G-LS, i.e., better segmentation accuracy. We observe that H-LS outperforms G-LS in 21 out of 25 cases.

For the sake of brevity in the ensuing discussion, the level set scheme using the proposed speed function (3) is referred as “H-LS” and the one using the conventional speed function (2) as “G-LS”. Results for both algorithms were evaluated against reference standard segmentations of positive remodeling wall regions that were provided by clinical specialists. Representative segmentation results are illustrated in Fig. 3 showing H-LS gives segmentations much closer to the reference positive remodeling regions, while G-LS tends to leak and propagate beyond the vessel wall boundary. We use the root mean squared distance (RMSD) of the segmented surface from the surface of the reference segmentation to quantify performance. We add the vessel lumen to both the reference and segmented volumes prior to calculating RMSD to mitigate large distances from false vessel wall detections opposite to the actual positive wall remodeling. Fig. 4 compares the RMSD values for G-LS and H-LS over twenty five datasets. The proposed H-LS method consistently outperforms standard G-LS. The mean RMSD over the twenty five cases was found to be 0.65 mm for G-LS and 0.47 mm for H-LS, giving a 28% improvement on average. Furthermore, a paired t-test between the RMSDs produced by H-LS and G-LS gives a p-value of 0.000032 suggesting statistically significant improvement.

6 Conclusion

We introduced a method for naturally incorporating nonparametric intensity priors of foreground and background objects into a local, nonparametric dissimilarity measure as a speed function term in the level set framework. The method improves segmentation accuracy for low contrast regions by providing a counterbalance to gradient-based front propagation. We demonstrated the method on a

database of vessels with positive wall remodeling. While vessel lumen narrowing or stenosis can be detected by directly analyzing the vessel lumen, our technique allows for the accurate segmentation of the remodeled wall region, a necessary step for quantifying pathology. Our approach is applicable to other low contrast segmentation problems where exemplars are available. The low-level segmentation scheme can also be viewed as an independent step or combined with additional analysis and classification approaches to aid in diagnosis.

References

1. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. *International Journal of Computer Vision* 1(4), 321–331 (1988)
2. Xu, C., Prince, J.L.: Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing* 7(3), 359–369 (1998)
3. Sethian, J.A.: A fast marching level set method for monotonically advancing fronts. In: *Proceedings of the National Academy of Sciences*, pp. 1591–1595 (1996)
4. Malladi, R., Sethian, J.A., Vemuri, B.C.: Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 17(2), 158–175 (1995)
5. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. *International Journal of Computer Vision* 22(1), 61–79 (1997)
6. Leventon, M., Faugeras, O., Grimson, W.: Level set based segmentation with intensity and curvature priors. In: *Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 4–11 (2000)
7. Paragios, N., Deriche, R.: Geodesic active regions and level set methods for supervised texture segmentation. *International Journal of Computer Vision* 46(3), 223–247 (2002)
8. Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics* 42(5), 577–685 (1989)
9. Chan, T., Vese, L.: Active contours without edges. *IEEE Transactions on Image Processing* 10(2), 266–277 (1999)
10. Baillard, C., Barillot, C.: Robust 3D segmentation of anatomical structures with level sets. In: *Medical Image Computing and Computer-assisted Intervention*, pp. 236–245 (2000)
11. Rath, Y., Michailovich, O., Malcolm, J., Tannenbaum, A.: Seeing the unseen: Segmenting with distributions. In: *International Conference on Signal and Image Processing* (2006)
12. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision* 40(2), 99–121 (2000)
13. Levina, E., Bickel, P.: The earth mover’s distance is the Mallows distance: some insights from statistics. In: *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 251–256 (2001)
14. Glagov, S., Weisenberg, E., Zarins, C.K., Stankunavicius, R., Kolettis, G.J.: Compensatory enlargement of human atherosclerotic coronary arteries. *New England Journal of Medicine* 316(22), 1371–1375 (1987)
15. Crawford, M., DiMarco, J.P., Paulus, W.J. (eds.): *Cardiology*. Elsevier, Amsterdam (2004)