IPMI 2001: 17th International Conference, pages 431–437 (preprint)

Nonlinear Edge Preserving Smoothing and Segmentation of 4-D Medical Images Via Scale-Space Fingerprint Analysis

Bryan W. Reutter[1,2], V. Ralph Algazi[2], and Ronald H. Huesman[1]

[1] Center for Functional Imaging, Lawrence Berkeley National Laboratory
University of California, Berkeley, CA 94720, USA
`http://cfi.lbl.gov/{~reutter, ~huesman}`
[2] Center for Image Processing and Integrated Computing
University of California, Davis, CA 95616, USA
`http://info.cipic.ucdavis.edu/~algazi`

Abstract. An approach is described which has the potential to unify edge preserving smoothing with segmentation based on differential edge detection at multiple scales. The analysis of n-D data is decomposed into independent 1-D problems. Smoothing in various directions along 1-D profiles through n-D data is driven by local structure separation, rather than by local contrast. Analytic expressions are obtained for the derivatives of the edge preserved 1-D profiles. Using these expressions, multidimensional edge detection operators such as the Laplacian or second directional derivative can be composed and used to segment n-D data. The smoothing and segmentation algorithms are applied to simulated 4-D medical images.

1 Introduction

Nonlinear edge preserving smoothing often is performed prior to medical image segmentation. The goal of the nonlinear smoothing is to improve the accuracy of the segmentation by preserving significant changes in image intensity, while smoothing random noise fluctuations. Methods include median filtering and gray-scale morphology [6], and spatially varying smoothing driven by local contrast measures [1] or nonlinear diffusion [8, 9]. By comparison, spatially invariant linear smoothing uniformly blurs boundaries in reducing noise, thus adversely affecting the accuracy of the subsequent segmentation.

Rather than irreversibly altering the data prior to segmentation, the approach described here has the potential to unify nonlinear edge preserving smoothing with segmentation based on differential edge detection at multiple scales. The analysis of multidimensional (n-D) image data is decomposed into independent 1-D problems that can be solved relatively quickly. Smoothing in various directions along 1-D profiles through n-D data is driven by a measure of local structure separation, rather than by a local contrast measure. The elementary 1-D smoothing algorithm is described in Section 2 and is generalized to arbitrary dimension in Section 3.

In addition, analytic expressions are obtained for the derivatives of the edge preserved 1-D profiles. Using these expressions and the methods described in Section 3, multidimensional edge detection operators such as the Laplacian or the second derivative in the direction of the image intensity gradient can be composed and used to segment n-D data. Computer simulations are used in Section 4 to evaluate the performance of 4-D versions of the n-D smoothing and segmentation algorithms. Preliminary results of a 3-D version of the n-D smoothing algorithm were presented in [2].

Potential applications of these methods include 4-D spatiotemporal segmentation of respiratory gated cardiac positron emission tomography (PET) transmission images to improve the accuracy of attenuation correction [4], and 4-D spatiotemporal segmentation of dynamic cardiac single photon emission computed tomography (SPECT) images to facilitate unbiased estimation of time activity curves and kinetic parameters for left ventricular volumes of interest [3].

2 1-D Recursive Multiscale Blending

Given linearly smoothed versions of a 1-D signal $\mathrm{f}(x)$ and its first two derivatives at J scales, one can perform nonlinear edge preserving smoothing as follows. The linearly smoothed versions of $\mathrm{f}(x)$ are denoted by $\bar{\mathrm{f}}(x, a_j)$, and the linearly smoothed first and second derivatives are denoted by $\bar{\mathrm{f}}^{(1)}(x, a_j)$ and $\bar{\mathrm{f}}^{(2)}(x, a_j)$, respectively, for $j = 1, \ldots, J$. The scale coordinate a controls the width of the convolution kernels used in the linear filtering. The kernels are based on the uniform cubic B-spline basis function and its first two derivatives [7]. The cubic B-spline has a support of $4a$ and approximates a Gaussian with a standard deviation, σ, of $\sqrt{1/3}\,a$. Dyadic sampling of the scale coordinate a is used, yielding $a_j = 2^{j-1}a_1$.

The nonlinearly smoothed versions of $\mathrm{f}(x)$, denoted by $\tilde{\mathrm{f}}(x, a_j)$, are obtained by recursively blending the linearly smoothed versions:

$$\tilde{\mathrm{f}}(x, a_j) = \begin{cases} \bar{\mathrm{f}}(x, a_1) & j = 1 \\ [1 - \mathrm{C}_j(x)]\,\tilde{\mathrm{f}}(x, a_{j-1}) + \mathrm{C}_j(x)\bar{\mathrm{f}}(x, a_j) & j = 2, \ldots, J. \end{cases} \tag{1}$$

The blending functions $\{\mathrm{C}_j(x);\ j = 2, \ldots, J\}$ are constrained to range between zero and one and play a role similar to that of the spatially varying diffusion coefficients used in typical implementations of edge preserving smoothing via nonlinear diffusion (e.g., [8, 9]). When $\mathrm{C}_j(x_0) = 0$, smoothing stops in the neighborhood of x_0 and $\tilde{\mathrm{f}}(x_0, a_j)$ remains unchanged from the value $\tilde{\mathrm{f}}(x_0, a_{j-1})$ obtained using nonlinear smoothing at the previous, finer scale. When $\mathrm{C}_j(x_0) = 1$, smoothing is unabated and $\tilde{\mathrm{f}}(x_0, a_j)$ is set to the value $\bar{\mathrm{f}}(x_0, a_j)$ obtained using linear smoothing at the current, coarser scale. Although the recursive multiscale blending cannot be characterized as nonlinear diffusion, it shares the desirable property of generating no spurious extrema, in the following sense. It can be shown that the nonlinearly smoothed signal $\tilde{\mathrm{f}}(x, a_j)$ is a convex combination of the linearly smoothed signals $\{\bar{\mathrm{f}}(x, a_i);\ i = 1, \ldots, j\}$ for all x, and therefore is bounded by the extrema of the linearly smoothed signals.

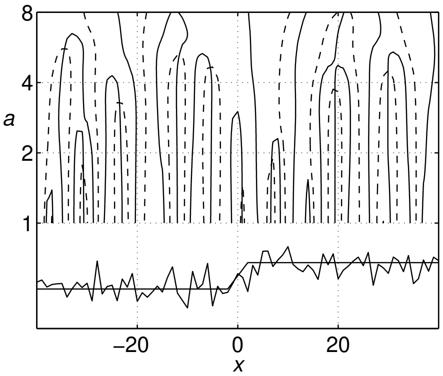

Fig. 1. Augmented scale-space fingerprint for a noisy ramp edge of width four and a contrast to noise ratio of 2.5. Solid fingerprint lines depict the zero-crossing locations of $\bar{\mathrm{f}}^{(2)}(x,a)$ (i.e., edge and ledge locations) over a continuum of scales. Dashed lines depict the zero-crossing locations of $\bar{\mathrm{f}}^{(1)}(x,a)$ (i.e., ridge and trough locations). Below the fingerprint, the noiseless edge is shown with the noisy edge.

The multiscale blending functions $\{\mathrm{C}_j(x);\, j = 2,\ldots,J\}$ are defined via the following analysis (presented in more detail in [2]) of the augmented scale-space fingerprint for $\mathrm{f}(x)$. The augmented scale-space fingerprint is a graphical depiction of the locations of the zero-crossings of the first two derivatives of the linearly smoothed signal as a function of scale (Fig. 1). At a particular scale a_j, each zero-crossing location of $\bar{\mathrm{f}}^{(2)}(x,a_j)$ is labeled as either a local maximum (edge) or local minimum (ledge) in gradient magnitude, depending on its proximity to nearby zero-crossing locations of $\bar{\mathrm{f}}^{(1)}(x,a_j)$ (i.e., ridges and troughs). For each edge location, the distance separating the ridge, trough, or ledge on either side of the edge is calculated. The blending function $\mathrm{C}_j(x)$ is then assigned a value ranging between zero and one at the edge location, based on the separation distance and the heuristic that larger separation distances are mapped to smaller blending function values. $\mathrm{C}_j(x)$ is then defined for all x by interpolating the values at the edge locations with a piecewise quartic spline whose first through third derivatives are zero at the edge locations.

3 *n*-D Smoothing and Segmentation

Edges can be preserved in n-D data by applying the 1-D smoothing algorithm described in Section 2 independently along the coordinate axis directions, as well as along the diagonal directions of the 2-D planes spanned by the coordinate axes, and averaging the results. This will be referred to as multidirectional 1-D processing, and builds on the work described in [9], in which processing was performed only along the coordinate axis directions. The information obtained along the diagonal directions allows the characterization of the first and second order differential properties of the data in any direction. Using this additional information, multidimensional edge detection operators such as the Laplacian or the second derivative in the direction of the image intensity gradient can be composed and used to segment the data as follows.

The n-D data array is denoted by $\mathrm{f}(\mathbf{x})$, where $\mathbf{x} = [\,x_1 \cdots x_n\,]^{\mathrm{T}}$ is the position vector for the domain of the data and "$[\;]^{\mathrm{T}}$" denotes the matrix transpose. The

1-D profile passing through the point $\mathbf{x}_0$ in the direction $\mathbf{v}_0$ is denoted by

$$\mathrm{f}_{\mathbf{x}_0,\mathbf{v}_0}(s) = \mathrm{f}(\mathbf{x}_0 + s\mathbf{v}_0), \tag{2}$$

where $\mathbf{v} = [v_1 \cdots v_n]^{\mathrm{T}}$ is a unit vector and s is an arc length parameter. The relationships between the first and second derivatives of $\mathrm{f}_{\mathbf{x},\mathbf{v}}(s)$ and the first and second order partial derivatives of the n-D data $\mathrm{f}(\mathbf{x})$ are

$$\frac{d\mathrm{f}_{\mathbf{x},\mathbf{v}}}{ds} = \mathbf{v} \cdot \nabla \mathrm{f} = \mathbf{v}^{\mathrm{T}}\mathbf{g} \qquad \frac{d^2\mathrm{f}_{\mathbf{x},\mathbf{v}}}{ds^2} = \mathbf{v} \cdot \nabla[\mathbf{v} \cdot \nabla \mathrm{f}] = \mathbf{v}^{\mathrm{T}}\mathbf{H}\mathbf{v}, \tag{3}$$

where $\mathbf{g}(\mathbf{x})$ is the gradient vector and $\mathbf{H}(\mathbf{x})$ is the Hessian matrix. One can write $\mathbf{v}^{\mathrm{T}}\mathbf{H}\mathbf{v}$ as the inner product $\mathbf{w}^{\mathrm{T}}\mathbf{h}$ of the $(\frac{n^2+n}{2})$-element vectors

$$\mathbf{w} = \left[v_1^2 \;\, 2v_1v_2 \;\, \cdots \;\, 2v_1v_n \;\, v_2^2 \;\, 2v_2v_3 \;\, \cdots \;\, 2v_2v_n \;\, \cdots \;\, v_{n-1}^2 \;\, 2v_{n-1}v_n \;\, v_n^2\right]^{\mathrm{T}} \tag{4}$$

$$\mathbf{h} = \left[H_{11} \; H_{12} \cdots H_{1n} \; H_{22} \; H_{23} \cdots H_{2n} \cdots H_{(n-1)(n-1)} \; H_{(n-1)n} \; H_{nn}\right]^{\mathrm{T}}, \tag{5}$$

where $H_{ij} = \frac{\partial^2 \mathrm{f}}{\partial x_i \partial x_j}$. Thus, given derivative estimates in all 1-D profiles along the coordinate axis directions and the diagonal directions of the 2-D planes spanned by the coordinate axes (for a total of n^2 directions), one can compute least squares estimates of the gradient vector $\mathbf{g}(\mathbf{x})$ and the vector $\mathbf{h}(\mathbf{x})$ of Hessian matrix elements as follows. The n^2 direction vectors for the 1-D profiles and the corresponding $\mathbf{w}$ vectors are stored in the matrices

$$\mathbf{V} = \left[\mathbf{v}_1 \cdots \mathbf{v}_{n^2}\right]^{\mathrm{T}} \qquad \mathbf{W} = \left[\mathbf{w}_1 \cdots \mathbf{w}_{n^2}\right]^{\mathrm{T}}. \tag{6}$$

The first and second derivatives along the 1-D profiles are stored in the vectors

$$\mathbf{f}^{(1)}(\mathbf{x}) = \left[\frac{d\mathrm{f}_{\mathbf{x},\mathbf{v}_1}}{ds} \cdots \frac{d\mathrm{f}_{\mathbf{x},\mathbf{v}_{n^2}}}{ds}\right]^{\mathrm{T}} \qquad \mathbf{f}^{(2)}(\mathbf{x}) = \left[\frac{d^2\mathrm{f}_{\mathbf{x},\mathbf{v}_1}}{ds^2} \cdots \frac{d^2\mathrm{f}_{\mathbf{x},\mathbf{v}_{n^2}}}{ds^2}\right]^{\mathrm{T}}. \tag{7}$$

It can be shown that the unweighted least squares estimates for the gradient vector $\mathbf{g}(\mathbf{x})$ and the vector $\mathbf{h}(\mathbf{x})$ of Hessian matrix elements are

$$\hat{\mathbf{g}}(\mathbf{x}) = \left[\mathbf{V}^{\mathrm{T}}\mathbf{V}\right]^{-1}\mathbf{V}^{\mathrm{T}}\mathbf{f}^{(1)} \qquad \hat{\mathbf{h}}(\mathbf{x}) = \left[\mathbf{W}^{\mathrm{T}}\mathbf{W}\right]^{-1}\mathbf{W}^{\mathrm{T}}\mathbf{f}^{(2)}. \tag{8}$$

Using these estimates, one can compose multidimensional edge detection operators such as the Laplacian, trace($\hat{\mathbf{H}}$), or the second derivative in the direction of the gradient, weighted by the squared magnitude of the gradient, $\hat{\mathbf{g}}^{\mathrm{T}}\hat{\mathbf{H}}\hat{\mathbf{g}}$.

4 4-D Smoothing and Segmentation Simulations

A 4-D version of the n-D smoothing algorithm was applied to simulated respiratory gated PET transmission images generated using the Mathematical Cardiac Torso (MCAT) phantom [5]. The 4-D image array was composed of 40 contiguous 5 mm-thick transverse slices at 15 respiratory phases. Each transverse slice had 80×80 pixels with pixel size 5×5 mm. Diaphragm and heart motion of 15 mm in the superior-inferior direction was simulated, in conjunction with chest wall diameter changes of 9.8 mm in the left-right direction and 20 mm in the anterior-posterior direction. Gaussian white noise was added to the images to yield contrast to noise ratios of 5.0 at the air-soft tissue boundary and 3.5 at the soft tissue-lung boundaries (Fig. 2a).

The 1-D smoothing algorithm was applied independently along the x, y, z, and t axes of the noisy 80×80×40×15 dataset, as well as along the 12 diagonal directions of the 2-D planes spanned by the axes. Multiscale linear 1-D filtering was performed in each of the 16 directions using a pre-smoother followed by cubic B-spline-based smoothing and differentiation operators operating at three different scales. The filters combined to yield kernels with supports 1×7, 1×11, and 1×19, which approximated Gaussians with $\sigma = 1$, $\sqrt{2}$, and $\sqrt{6}$ pixels, respectively. For comparison, linear smoothing was also performed using a 5×5×5×5 separable filter, which approximated a 4-D Gaussian with $\sigma = 0.70$ pixels. This small scale separable filter was designed to yield the same noise reduction for independent, identically distributed Gaussian noise, as that obtained by averaging the outputs of the 16 large scale (1×19 B-spline-based) linear 1-D smoothing filters (Figs. 2c,e). Fig. 2g shows the result of averaging the outputs of the 16 large scale nonlinear 1-D smoothing filters obtained using recursive multiscale blending. The differences between the results are subtle. The large scale nonlinear multidirectional 1-D filter and the small scale separable filter blurred the edges the least, while the large scale linear multidirectional 1-D filter blurred the edges the most (Fig. 2b). The linear and nonlinear multidirectional 1-D smoothing results were obtained using an average of 5.8 minutes of processing for each of the 16 directions (195 MHz R10000-based SGI workstation).

Results of segmenting the images using 4-D second directional derivative operators are shown in Figs. 2d,f,h. For the linear and nonlinear multidirectional 1-D processing, the 4-D gradient vector and Hessian matrix were calculated in 17 minutes using the methods described in Section 3. For respiratory phase 8, 3-D models for the second directional derivative zero-crossing surfaces were constructed in less than one minute using the methods described in [4]. The large scale nonlinear multidirectional 1-D operator and the small scale separable operator yielded comparable segmentations. Relatively accurate lung surface models were constructed, to which were attached spurious surface elements. For the large scale linear multidirectional 1-D operator, there were fewer spurious surface elements and the lung surface models were less accurate.

5 Future Directions

The computer simulations in Section 4 demonstrate that nonlinear edge preserving smoothing and segmentation of 4-D medical images can be performed in a timely manner on a workstation. Unlike typical implementations based on nonlinear diffusion, recursive multiscale blending requires only a small, fixed number (3–5) of iterations. Although performed serially here, the computations

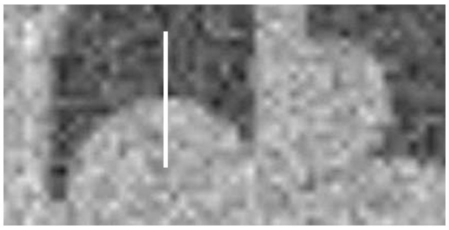

(a) original noisy image

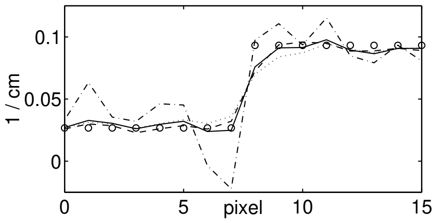

(b) edge at diaphragm

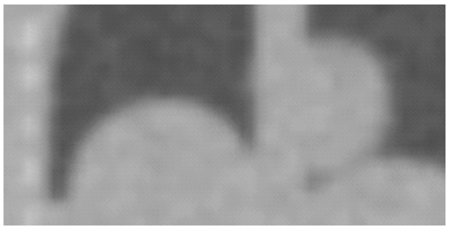

(c) 5×5×5×5 linear smoothing

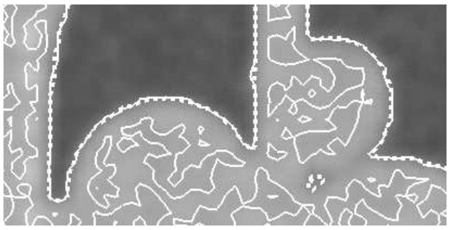

(d) 5×5×5×5 linear smoothing

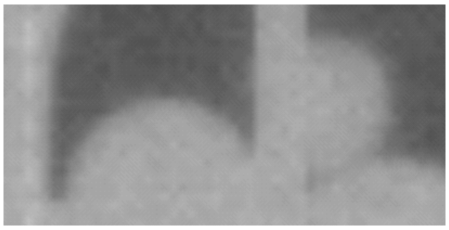

(e) 16×(1×19) linear smoothing

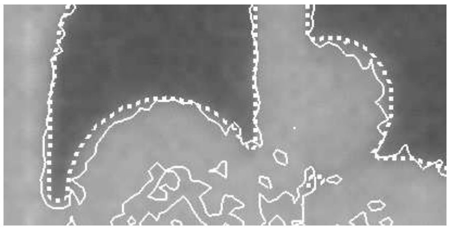

(f) 16×(1×19) linear smoothing

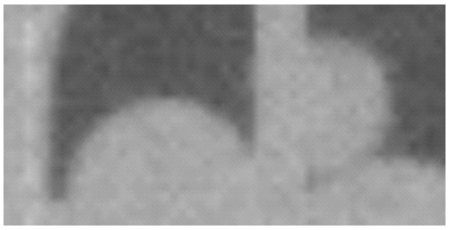

(g) 16×(1×19) nonlinear smoothing

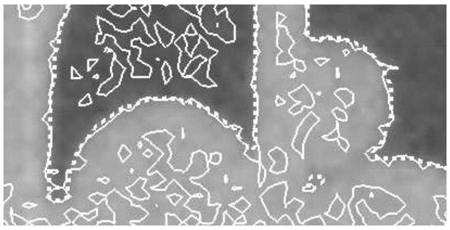

(h) 16×(1×19) nonlinear smoothing

Fig. 2. Smoothing and segmenting simulated 4-D respiratory gated PET transmission images. (a) Noisy 52×26 pixel sub-image from a coronal cross section. The right dome of the diaphragm is the larger, semicircular structure on the left. The heart is the smaller, circular structure on the right. (b) Profile through right dome of diaphragm, depicted by the white segment in (a). The circles and the dot-dashed line depict noiseless and noisy simulated values, respectively. The dashed, dotted, and solid lines depict values obtained by (c) small scale separable, (e) large scale linear multidirectional 1-D, and (g) large scale nonlinear multidirectional 1-D filtering, respectively. (d,f,h) Segmentation results for (c,e,g), respectively, are depicted as solid lines. The dotted lines depict the true soft tissue-lung boundaries.

can be massively parallelized. Additional work is needed to optimize the multiscale blending functions with respect to spurious zero-crossings in the derivatives of the nonlinearly smoothed data. With the goal of improving the preservation of fine details, further investigation is needed to perform weighted least squares estimation of a 4-D dataset and its partial derivatives from the results of performing recursive multiscale blending in multiple directions.

Acknowledgments

The authors thank the University of North Carolina Medical Imaging Research Laboratory for making the MCAT phantom available.

This work was supported by the National Heart, Lung, and Blood Institute of the US Department of Health and Human Services under grant P01-HL25840; by the Director, Office of Science, Office of Biological and Environmental Research, Medical Sciences Division of the US Department of Energy under contract DE-AC03-76SF00098; and by the University of California MICRO program. This work was developed in part using the resources at the US Department of Energy National Energy Research Scientific Computing (NERSC) Center.

References

1. Kitamura, K., Iida, H., Shidahara, M., Miura, S., Kanno, I.: Noise reduction in PET attenuation correction using non-linear Gaussian filters. IEEE Trans. Nucl. Sci., **47** (2000) 994–999
2. Reutter, B.W., Algazi, V.R., Huesman, R.H.: Computationally efficient nonlinear edge preserving smoothing of n-D medical images via scale-space fingerprint analysis. In Ulma, M. (ed.), 2000 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (2001, in press)
3. Reutter, B.W., Gullberg, G.T., Huesman, R.H.: Direct least-squares estimation of spatiotemporal distributions from dynamic SPECT projections using a spatial segmentation and temporal B-splines. IEEE Trans. Med. Imag., **19** (2000) 434–450
4. Reutter, B.W., Klein, G.J., Huesman, R.H.: Automated 3-D segmentation of respiratory-gated PET transmission images. IEEE Trans. Nucl. Sci., **44** (1997) 2473–2476
5. Segars, W.P., Lalush, D.S., Tsui, B.M.W.: Modeling respiratory mechanics in the MCAT and spline-based MCAT phantoms. In Seibert, J.A. (ed.), 1999 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (2000) 985–989
6. Sternberg, S.R.: Grayscale morphology. Comput. Vis. Graph. Image Proc., **35** (1986) 333–355
7. Wang, Y.-P., Lee, S.L.: Scale-space derived from B-splines. IEEE Trans. Patt. Anal. Mach. Intell., **20** (1998) 1040–1055
8. Weickert, J.: A review of nonlinear diffusion filtering. In ter Haar Romeny, B., Florack, L., Koenderink, J., and Viergever, M. (eds.), Scale-Space Theory in Computer Vision: Proceedings of the First International Conference (1997) 3–28
9. Weickert, J., ter Haar Romeny, B.M., Viergever, M.A.: Efficient and reliable scheme for nonlinear diffusion filtering. IEEE Trans. Image Proc., **7** (1998) 398–410