Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Fast Edge-Preserving Filtering for 3D+t Echocardiographic Volume Rendering

Oscar Yanez-Suarez and Jean François Lerallut

Abstract— **A novel edge-preserving filtering method based on unsupervised clustering within the joint four-dimensional location-intensity space of volume echocardiographic data is presented. A vector codebook, formed by vectors that lie close to the modes of the joint data distribution is derived through an iterative vector quantization procedure. Distribution estimates are non-parametric, kernel based, and thus the quality of the quantization can be controlled by adequate choice of the kernel parameters.**

I. Introduction

Three-dimensional (3D) echocardiography has gained its prominent place in cardiologic assessment given that it provides adequate data for improved left ventricular quantification and better determination of anatomical relationships. Moreover, modern 3D+t echocardiographic technologies such as 2D crystal arrays [1] allow, for example, the accurate tracking of heart movements [2]. However, the 3D+t echocardiographic imaging studies are composed of a high volume of data, typically in the order of 10-20 megavoxels per frame, with 15-30 frames per cardiac cycle. While dedicated hardware for real-time volume rendering of the image sets is readily available [3], the embedded image processing is usually limited to operator-dependent manual thresholding. Recent contributions aim at providing better tissue segmentation, which is a necessary step for left ventricular analysis or adequate opacity assignment in volume rendering techniques [4], [5].

Besides the large amount of image data, inherent characteristics of 3D ultrasound impose difficulties for reliable segmentation. We present a robust method for edge-preserving filtering of 3D ultrasound frames based on vector quantization over the joint range-location space, which is achieved by optimizing a Kullback-Liebler divergence of the non-parametric probability density estimate of the volume data. Vector quantization in the joint space provides echocardiographic noise smoothing without degradation of relevant edges while producing a segmented scene that is ready for labeling prior to opacity assignment in a volume rendering context.

II. Methodology

The computation of the vector quantization codebook and its utilization for edge-preserving filtering are detailed in the following subsections.

This work was supported by a research grant from LAFMI, the Franco-Mexican Laboratory of Informatics

O. Yanez is with the Neuroimaging Laboratory, Department of Electrical Engineering, Universidad Autonoma Metropolitana - Iztapalapa, Mexico email: yaso@xanum.uam.mx

J.F. Lerallut and O. Yanez are with the Image Processing Laboratory, UMR6600, Université de Technologie de Compiègne, France. email: jflerallut@utc.fr

A. Density Estimation in the Joint Range-Location Space

One frame of 3D+t echocardiographic data, recorded at instant t, is an $N \times M \times P$ volume $\mathbf{V}_t(x, y, z)$ of echo intensities or *ranges*. The joint *range-location* representation of the frame is constructed by building the four-dimensional vectors:

$$\mathbf{J}_{k,t} = \left[\begin{array}{cccc} x_k & y_k & z_k & \mathbf{V}_t(x_k, y_k, z_k) \end{array} \right]' \tag{1}$$

where k varies from 0 to $K-1, K < N \times M \times P$, to include all the information-bearing voxels of the frame (given the geometry of the probe and the beam parameters, the geometry of the imaging cone within $\mathbf{V}_t$ can be determined).

Echo data in the joint space is actually a realization of a vector random process with unknown probability density $f(\mathbf{J}_{k,t})$. A non-parametric estimate of this density at a given point $\mathbf{x}$ in the joint space can be constructed from the $\{\mathbf{J}_{k,t}\}$ set by means of a kernel estimate:

$$\hat{f}(\mathbf{x}) = \frac{1}{Kh^4} \sum_{k=0}^{K-1} e^{-\frac{\|\mathbf{x}-\mathbf{J}_{k,t}\|^2}{h}} \tag{2}$$

where h is the smoothing parameter of the density estimate [6]. For reasons that will become aparent in the next section, the density gradient is also a relevant quantity to estimate. An estimate of the gradient of the density function can be written in terms of the estimate in (2) by direct differentiation:

$$\hat{\nabla} f(\mathbf{x}) = \frac{1}{Kh^4} \sum_{k=0}^{K-1} \nabla e^{-\frac{\|\mathbf{x}-\mathbf{J}_{k,t}\|^2}{h}} \tag{3}$$

B. Quality of the Density Estimates

A measure of the similarity of the density estimate in (2) to the true (but unknown) density $f(\mathbf{x})$, can be established in terms of the Kullback-Leibler divergence [7]:

$$I_{\hat{f}\|f} = -\int f(\mathbf{x}) \log \frac{\hat{f}(\mathbf{x})}{f(\mathbf{x})} \, d\mathbf{x} \tag{4}$$

which will be zero if the estimated and the real densities are the same. Therefore, those values of $\mathbf{x}$ where the divergence is minimized are the points in joint space where the density is best estimated by (2). To minimize $I_{\hat{f}\|f}$ with respect to $\mathbf{x}$ it is required that:

$$\nabla I_{\hat{f}\|f} = -\int \left(\left[\frac{1}{\hat{f}(\mathbf{x})} \hat{\nabla} f(\mathbf{x}) \right] f(\mathbf{x})\, d\mathbf{x} \right) |_{\mathbf{x}=\mathbf{w}_k} \equiv 0 \quad (5)$$

where use has been made of the density gradient estimate in (3) and $\mathbf{w}_k$ is the solution point for minimum divergence. Solving for $\mathbf{w}_k$ is not directly possible because the true density is unknown. However, it is seen from equation (5) that

$$\nabla I_{\hat{f}\|f} = -E\left\{ \frac{1}{\hat{f}(\mathbf{x})} \hat{\nabla} f(\mathbf{x}) \right\} \quad (6)$$

where $E\{\cdot\}$ denotes the expectation operator. Thus it is possible to write the Robbins-Monro stochastic approximation iterations [8] for solving (5) as:

$$\mathbf{w}_k^{t+1} = \alpha_t \frac{\hat{\nabla} f(\mathbf{w}_k^t)}{\hat{f}(\mathbf{w}_k^t)} + \mathbf{w}_k^t \quad (7)$$

where α_t, the learning rate, *must* be a decreasing function of the iteration number t.

Careful observation of equation (7) shows that such iterations are actually performing gradient ascent over the estimated probability density, and therefore are *mode-seeking* iterations, that is, when starting from an initial point $\mathbf{w}_k^0$ in the joint space, the iterations will take that point towards the closest local mode of the range-location density.

It is worth noting that the magnitude of the ascent step in each iteration is adaptively driven by the data through the concourse of three inter-compensating factors:

1) The estimate of the density gradient, which should be small near the modes of the distribution,
2) The density estimate itself, that should exhibit large values near the modes, and
3) The learning rate, which slows down the ascent as the iterations progress

Thus, iterations are self-limiting and will eventually stop.

C. Vector Quantization through Mode-Seeking

Making $\mathbf{w}_k^0 = \mathbf{J}_{k,t}, k = 0, \ldots, K-1$, that is, starting from the data in the volume image frame (1), the evaluation of iterations (7) will take each original voxel to its local range-location mode. It is reasonable to expect that voxels with similar intensities within a spatial neighborhood would converge towards the same local mode $\mathbf{J}_{m,t}$, which can be regarded as the *codeword* vector for the neighborhood. At the end of the iterations for all the voxels, vector quantization is achieved, producing a segmented volume.

Morever, substituting the original intensity of the voxel with the intensity of its associated mode (that is, replacing $\mathbf{V}_t(x_m, y_m, z_m)$ for $\mathbf{V}_t(x_k, y_k, z_k)$ within $\mathbf{J}_{k,t}$), a local nonlinear transformation of the original intensities is performed. Since voxels on either side of a 3D edge will cluster in different range-location modes, this substitution will preserve edges. Moreover, since noisy voxels will not cluster with those voxels from the underlying homogeneous volume region, noise filtering will also be achieved.

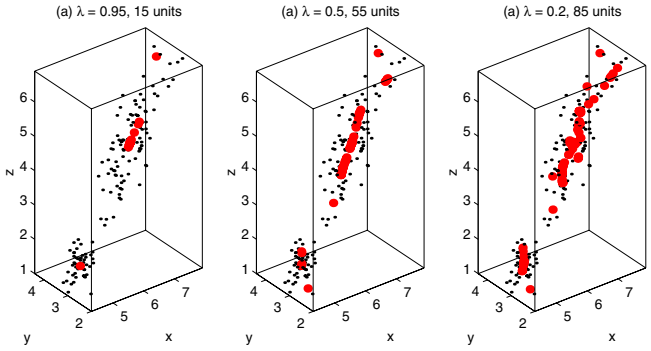

Fig. 1. Effect of varying the learning rate schedules. Original data (black dots) correspond to the first three features of the well-known Iris data set. Red dots correspond to the locations of the trained codewords or *units*. All the schedules are exponential, λ^t, with the parameter values shown.

The edge-preserving filtering and segmentation procedure just described can be summarized as follows:

1) Set the initial codeword vectors to be equal to the training data set, i.e. $\mathbf{w}_k^0 = \mathbf{J}_{k,t}, k = 0, \ldots, K-1$
2) Select the kernel bandwidth. A possible value is proposed in [6]: $h = h_s \times 1.06\, std(\mathbf{x}) K^{-0.2}$, with $std(\mathbf{x})$ equal to the standard deviation of the training data set
3) Select an initial learning rate. The only requirement is that the learning rate schedule must be strictly decreasing with the iteration number. An exponential rate, $\alpha_t = \lambda_0^t$, is a reasonable alternative
4) Update the codeword vectors using the Robbins-Monro iterations (7)
5) Consolidate those codeword vectors that are closer to each other than a given threshold
6) Repeat from step 4 until convergence

III. Results

A. Simulations

With the purpose of illustration of the vector quantization procedure, three features from the well-known Iris data set [9] are used in the experiments of Fig. 1. Starting from the original data (black dots), iterations (7) evolve into the subset of local modes shown in red. The learning rate in each case has an exponential schedule, $\lambda^t, \lambda < 1$. Larger λs favor the ascent of the codewords all the way to the modes, thus producing a small amount of final codeword vectors. On the other hand, smaller λs freeze the codewords faster, thus generating a larger codebook. Further, kernel bandwidth determines the degree of locality of the clusters.

Fig. 2 shows another straightforward application for color quantization. Since the original micrograph has 4096 gray levels, the joint range-location space is 3D. Coarse quantization

in this space produces a 5 tone image, with evident color distortion but nicely preserved shapes. Finer quantization yields a 114 gray level image that is visually undistinguishable from the original. In these examples, after codebook convergence, each pixel intensity is replaced by the intensity of the pixel's associated mode.

B. 3D+t Volume Rendering

Filtering/segmenting a frame of 3D echo requires vector quantization in 4D, over a huge amount of data points. As an approximation strategy, the computation of density and density gradient estimates, equations (2) and (3) respectively, can be performed only once over the original data grid, leaving the computation of function values at unseen locations to a local interpolation of the already computed points, and thus speeding up the entire processing. An implementation of the proposed scheme in VTK [10] takes approximately 10 seconds in standard PC hardware (1.6 GHz Pentium IV, 2 GB RAM, Linux) to process an entire frame. Fig. 3 corresponds to a screenshot of our visualization program showing the volume rendered cavities after filtering and segmentation.

IV. Conclusions

A novel scheme for edge-preserving filtering and segmentation of 3D echocardiographic image sets has been introduced. Based on the optimization of the Kullback-Leibler divergence for non-parametric estimates of the joint range-location probability density of the image dataset, the procedure is fast and has a fixed computation time per volume frame. Process properties are the consequence of local vector quantization and the replacement of voxel intensities with those corresponding to the local modes of the joint distribution. Kernel bandwidth and learning rate schedule are free parameters that have to be tuned by the system user: stronger image smoothing requires large kernel bandwidths and slower learning rates.

References

[1] L. Sugeng, L. Weinert, K. Thiele, and R. M. Lang, "Real-time three-dimensional echocardiography using a novel matrix array transducer," *Echocardiography*, vol. 20, no. 7, pp. 623–635, 2003.

[2] Y. Suematsu, G. R. Marx, J. A. Stoll, and P. E. DuPont *et al*, "Three-dimensional echocardiographyguided beating-heart surgery without cardiopulmonary bypass: A feasibility study," *Journal of Thoracic and Cardiovascular Surgery*, vol. 128, pp. 579–587, 2004.

[3] Philips Electronics. (2006, March) Sonos 7500. [Online]. Available: http://www.medical.philips.com/us/products/ultrasound/cardiology/

[4] R. Valdes-Cristerna, J. R. Jimenez-Alanis, O. Yanez-Suarez, J. F. Lerallut, and V. Medina, "Texture-based echocardiographic segmentation using a non-parametric estimator and an active contour model," in *Proceedings of the 26th Annual International Conference of the IEEE EMBS*. IEEE Press, 2004, pp. 1806–1809.

[5] V. Zagrodsky and R. Shekhar, "Volume rendering of real-time 3D echocardiographic data," in *IEEE Visualization 2005 - (VIS'05)*. IEEE Press, 2005, p. 114.

[6] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. Chapman-Hall / CRC, 1986.

[7] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, 1991.

[8] H. Robbins and S. Monro, "A stochastic approximation method," *Annals of Mathematical Statistics*, vol. 22, pp. 400–407, 1951.

[9] D. J. Newman, S. Hettich, C. L. Blake, and C. J. Merz. (2006, March) UCI repository of machine learning databases. [Online]. Available: http://www.ics.uci.edu/∼mlearn/MLRepository.html

[10] Kitware Inc. (2006, March) VTK, visualization toolkit. [Online]. Available: http://public.kitware.com/VTK/

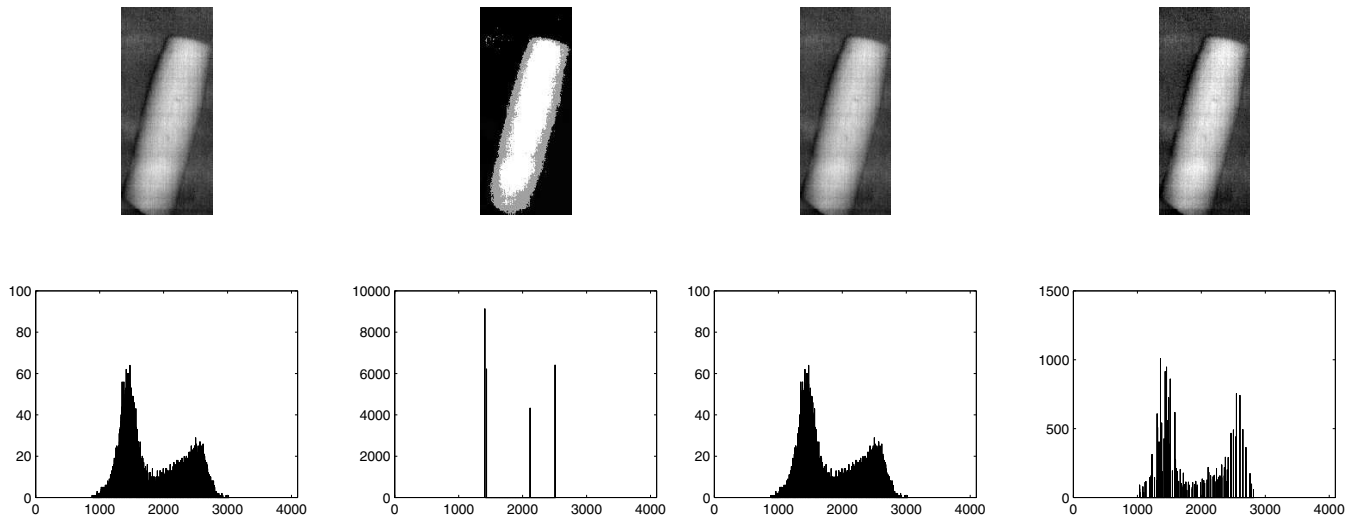

Fig. 2. Multiscale color quantization. Top, left to right: Original image at 12 bits per pixel (4096 gray levels); choice of a large bandwidth and slow learning rate produces a coarse grayscale quantization down to only five gray levels; orginal image, repeated for easy comparison; small bandwidth and fast learning schedule produces a more detailed result with 114 graytones. Bottom: corresponding histograms

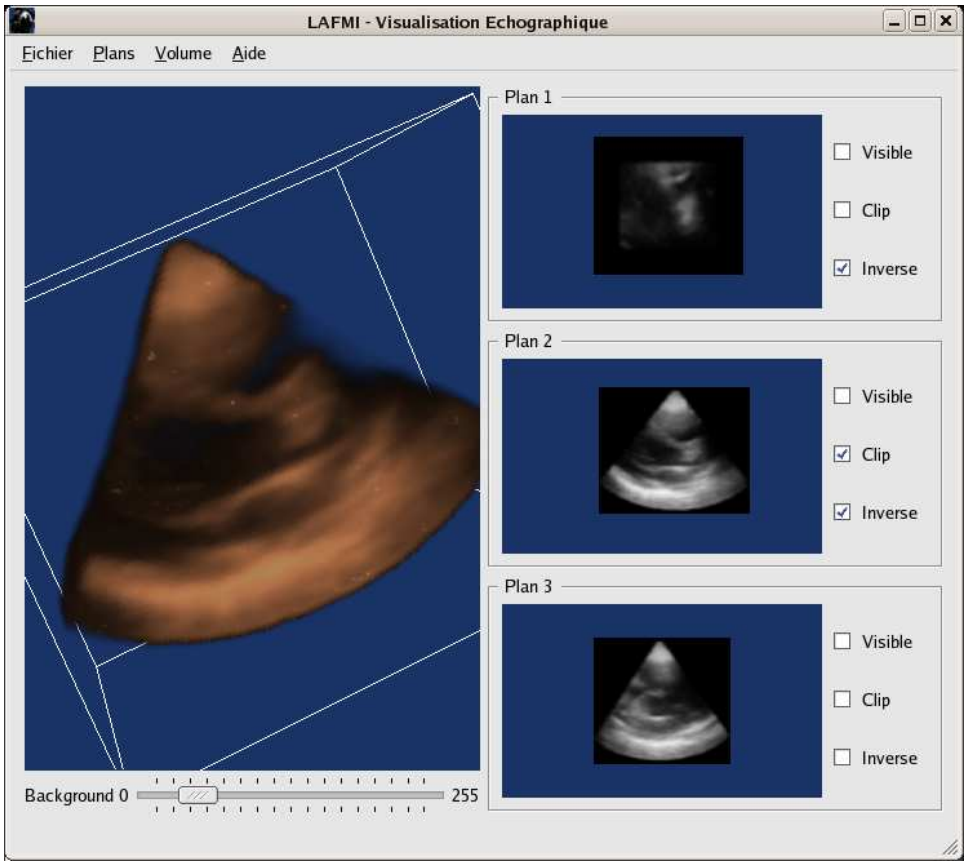

Fig. 3. Screenshot of 3D+t visualization after edge-preserving smoothing.