

Automatic heart localization from a 4D MRI dataset

Wolfgang Sörgel*

Telecommunications Institute
University of Erlangen-Nuremberg, Germany

Vincent Vaerman

Signal Processing Laboratory
Swiss Federal Institute of Technology, Lausanne, Switzerland

ABSTRACT

The purpose of the presented work is the automatic localization of the heart from 4D multi-slice magnetic resonance images (MRI). Well known active contour extraction techniques such as “snakes” or “balloons” require precise initialization which is mostly done interactively by the user in existing systems. A new method for the automatic initialization of such models is presented here for application on 4D MRI dataset acquired from the human heart. The method consists of two main steps: a global *localization* of the heart and a coarse *initialization* of the contours. Furthermore, it is shown how this initialization can be used for an automatic fine segmentation by an *active contour* model.

The temporal analysis of the heart beat cycle is well suited for *localization* purposes. A “temporal variance image” is thus first computed at each spatial slice location. These variance images consistently highlight the heart due to its wall movement and the heavy blood flow. By thresholding the variance images, projecting them into a single image, thresholding again and selecting the largest resulting object, a “binary confidence mask” is computed for the heart region. This mask allows us to extract one binary image of the heart for each spatial slice location, regardless of temporal location.

In the *initialization* stage, an “initial contour” is matched to each of these masked images by affine transform, adapting size, location, aspect ratio and orientation. Initial contours may be acquired from a predefined model. In absence of such a model, ellipses were successfully applied as generic initial contours. For this stage, 2D contours were used; however, extensions to 3D are straightforward.

The affine adapted contours are then considered as initialization for a *multi-step active contour model* for the accurate extraction of the heart walls: the contours are deformed according to the masked binary images, further refined on the temporal mean images for each spatial slice location, and finally the outer heart walls are tracked over time on the actual images at each spatial slice location. This approach, making use of existing active contour models yields an efficient and robust method for an exact extraction of the heart contours.

Keywords: cardiac imaging, MRI, 3D/4D image processing, contour extraction, segmentation

* At the time of this work, the author was with the Signal Processing Laboratory of the Swiss Federal Institute of Technology in Lausanne, Switzerland, within the European student exchange program ERASMUS.

1 INTRODUCTION

This paper presents a novel approach to heart localization in 4D magnetic resonance imaging (MRI) datasets. The provided dataset is composed of 17 parallel images of the heart and for each of these, 16 images representing the different phases of the heart cycle. The acquisition was cardiac-gated to provide the phase images. A spatial magnetic presaturation of the blood before it enters the image slice is applied. This leads to so-called “dark-blood” images with better contrasted contours. Some of these images are presented in Figure 1. Note that throughout this paper, the figures always display the third, seventh and eleventh slices –out of the 17-slices set– when different spatial slice location are presented.




Figure 1: Examples of transverse dark-blood images extracted from the 4D MRI dataset. (Top line: parallel spatial slice locations; bottom line: other temporal slice locations for corresponding slices)

Well known active contour model techniques –such as *snakes*⁶ or *balloons*⁵– provide proven methods for the fine segmentation of object contours. However these techniques need to be initialized sufficiently close to the target contour to guarantee stable convergence. This important problem leads to the task of finding an approximate position and pose of the heart within the given image set, and to adapt an initial contour model accordingly. By exploiting this kind of approach, one will be able to reuse existing accurate segmentation algorithms.

According to these different steps, the paper is organised as follows: first, a coarse but automatic localization of the heart is described in Section 2 leading to the definition of a mask representing the heart region; then an initial contour –a “template”– is fitted to the masked images in Section 3. This efficient initialization is then validated in Section 4 by applying it to active contour models. Finally, Section 5 draws the conclusions and gives some hints about future works.

2 AUTOMATIC LOCALIZATION OF THE HEART

In this section, a new approach is presented for the automatic localization of the heart in 4D image sets based on the temporal variance of the image gray levels. The method results in a binary image of the heart for each spatial slice location. These images are then used in Section 3 to adapt position, orientation and size of model templates which in turn are a suitable input for fine segmentation with active contour models or balloons.

Different approaches to the problem of localizing an object –in our case, the heart– in a cluttered scene could be used and combined, including:

- manual outlining of all contours by an operator;
- model-based approaches using spatial relation within the image scene such as relative position of the heart with respect to the outer upper body contour or other easily detectable organs as lungs or spine;
- region-based approaches using texture and gray level information;
- template matching, implemented by the generalized Hough transform¹;
- use of special properties of the tissue to be segmented.

As the goal is to develop a method that shall work without human intervention and for different image perspectives, the first two approaches are ruled out. Due to the large differences in gray level and texture for different MRI techniques, methods based on these features do not seem to be very promising either. Rigid template matching⁷ is computationally expensive and not well suited to detect a highly deformable object such as the heart, and has been confirmed experimentally. An example for a complex multi-step approach using fuzzy Hough transform¹¹ and region growing is reported by Philip *et al.*¹² However, the algorithm only works for transversal images and relies on the presence of other body features besides the heart in the images, such as chest wall and the left lung.

In this approach, the fact that the *heartbeat* motion is a unique and strong movement is exploited, independently of the perspective and field of view. The heart localization is based on the analysis of the image gray level variance over time. Details are presented in Section 2.1. The inclusion of spatial information leads to the definition of a mask in Section 2.2, while Section 2.3 briefly presents the direct adaptation of the method to datasets restricted to “2D + time” images.

2.1 Analysis of the temporal variance

This work is based on 4D image sets and subsets. A 4D image set is defined as a mapping $\mathbf{I}(\mathbf{p}) : \mathbb{I} \mapsto \mathbb{R}$ with the image set domain $\mathbb{I} = \{\mathbf{p} = (x, y, l, t)\}$ where $x = 1, 2, \dots, X$; $y = 1, 2, \dots, Y$ are the image coordinates, $l = 1, 2, \dots, L$ is the spatial slice location, and $t = 1, 2, \dots, T$ is the temporal slice location. This means the set contains $L \cdot T$ images $\mathbf{I}(\mathbf{x})$ with $\mathbf{x} = (x, y)$ at L different normalized spatial locations l orthogonal to the image coordinates (x, y) and T different normalized discrete times t . To simplify the notation throughout the paper, $\mathbf{I}_l^t(\mathbf{x})$ will denote a single image at spatial location l and temporal location –*e.g.* time– t .

2.1.1 Computation and Interpretation

Let \mathbf{I} be a given 4D image set. For analysis of the heart motion a set \mathbf{M} of L temporal *mean images* at each slice position l

$$\mathbf{M}_l = \frac{1}{T} \sum_{t=1}^T \mathbf{I}_l^t \quad (1)$$

is computed and from that a set \mathbf{V} of *variance images*

$$\mathbf{V}_l = \frac{1}{T-1} \sum_{t=1}^T (\mathbf{I}_l^t - \mathbf{M}_l)^2. \quad (2)$$

It was experimentally determined that within the outer contours of the heart, the temporal variance has large values with high probability. Outside the heart generally smaller values are found, independently of MR acquisition technique or perspective. The reason for this is the movement of the heart walls during the cardiac cycle and the associated heavy flow of blood in the heart. Figure 2 shows some examples for different spatial slice locations. Position, size and orientation of the heart can be well seen. Especially, the heart is well separated from the surrounding bone tissue, as opposed to the situation in actual images.

However it is also evident that there may be other areas of high variance in the image, caused for example by movement of the patient body or breathing during the image acquisition. Some movement artifacts will probably be reduced in the future by even faster acquisition techniques.³ One still needs to deal with variance peaks generated by other organs within the field of view, and with a reasonable amount of movement and imaging artifacts. Furthermore, the variance value is not constant within the heart: blood vessels and the walls of the left ventricle generally show higher variance. These facts are taken into account by binarization for the further processing, as discussed below.




Figure 2: Examples of computed variance images

An interesting side aspect of this work –aiming to localize and segment the outer heart contours– may be the direct use of the variance images for *diagnostic* purposes. Especially disturbed movement and blood flow in non-healthy hearts may be well visible. The idea has been discussed with some physicians: more experience with images of pathologic organs has to be gained before any final conclusions can be drawn.

2.1.2 Thresholding

The variance images \mathbf{V}_l are then thresholded to obtain binary images outlining the heart. Additional objects, resulting from the various noise sources described above, are still tolerated at this step. Experimentation showed that it is advantageous to threshold the \mathbf{V}_l so that a given number of pixels s_H –representing the *expected* maximal size of the heart– is set to one, and the rest to zero. This threshold s_H is first assumed to be known[†]. A completely data-driven automatic threshold estimation is suggested further in Section 2.1.3.

The threshold for binarization is found by first computing a histogram $h_{\mathbf{V}_l}(v)$ and then finding the largest

[†]This parameter can be easily estimated by a trained physician (in a metric unit needed then to be converted to pixels). Furthermore, the algorithm showed to be insensitive to miss-estimation of s_H up to a factor of 2.

value v_l^T for which the inequality

$$\int_{v_l^T}^{v_l^{\max}} h_{\mathbf{V}_l}(v)dv \leq s_H \quad (3)$$

holds. The thresholded variance images are then obtained by setting

$$\mathbf{V}_l^T(x, y) = \begin{cases} 1 & \text{if } \mathbf{V}_l^T(x, y) > v_l^T \\ 0 & \text{else} \end{cases} \quad (4)$$

Thus, approximately s_H pixels with the *largest variance values* are kept in each image. Note that v_l^T is generally different for each spatial slice location l .

Figure 3 shows some resulting images. It can be seen that the heart is well outlined in all images with only few blank spots within. On the other hand there are also noise clusters present which may in some cases be even larger than the heart. This is addressed by the use of a confidence mask described in Section 2.2. The results for different choices of s_H are compared in Figure 4 where the heart is visibly well outlined in all images. In the lower images, at a spatial slice location where the heart cross section is maximal, we note that some parts are missing for $s_H = 2500$. The algorithm will however work with both extreme cases shown in Figure 4, although they are close to the feasible limits.




Figure 3: Variance images after thresholding for different spatial slice locations

2.1.3 Automatic determination of the threshold

From Figure 4, it can be seen that the thresholded variance image contains only one large object for correctly chosen heart size. For underestimated heart size, the heart is broken up in several smaller parts; for larger objects, additional noise clusters are visible. This observation is used to suggest a method for the automatic optimal determination of s_H : the threshold $v_l^T(s_{H,i})$ are computed according to (3) for different values $s_{H,i}$ of the estimated heart size, resulting in different binary variance images for the different thresholds (see Figure 4). In each image, the size of all connected objects is computed and from that the ratio of the size of the largest object to the sum over all other objects.

The value $s_{H,\max}$ where this ratio is maximal has been found to be a good indicator of the heart size. As the latter processing is better suited to an overestimation, s_H was set to $s_H = \varepsilon n_{H,\max}$ with a multiplicative factor ε . In our experience, $1 \leq \varepsilon \leq 2$ were suitable, and the choice of $\varepsilon = 1.25$ was successfully used. A coarse iteration of s_H , using approximately ten different values, is enough for practical application. Furthermore, it is sufficient to determine s_H only for one slice and use the value found for all others as well.




Figure 4: Variance images thresholded to keep (from left to right) 2500, 5000 and 10000 pixels, for 2 (top and bottom) different spatial slice locations

2.1.4 Morphological cleaning of the binary images

The “salt-and-pepper” noise observed in the thresholded variance images is removed by a morphological cleaning operation.² This operation, although not absolutely necessary, was used to increase the accuracy and robustness of the process. For each pixel (x, y) in \mathbf{V}_l^T , let

$$\mathbf{V}_l^T(x, y) = \begin{cases} 1 & \text{if } \sum_{\Omega_r(x, y)} \mathbf{V}_l^T(u, v) > n_{\max} \\ 0 & \text{if } \sum_{\Omega_r(x, y)} \mathbf{V}_l^T(u, v) < n_{\min} \\ \mathbf{V}_l^T(x, y) & \text{else} \end{cases} \quad (5)$$

where $\Omega(x, y)$ is a neighborhood around pixel (x, y) . We used 5×5 square windows, yielding 24 neighbors for each pixel, $n_{\max} = 15$ and $n_{\min} = 11$. This parameter choice removes small objects of up to 10 pixels, erodes small protrusions and closes holes of up to 9 pixels. Note that the images in Figures 3 and 4 are already cleaned using (5).

2.2 Including spatial information: using a mask

It was shown in Section 2.1.2 that the heart region is visible in all L slices of the thresholded variance images \mathbf{V}_l^T , whereas the noise is more randomly distributed within the images, and other organs typically extend only over few slices. This is used in the following steps to build a binary mask outlining the heart region:

1. The thresholded and cleaned variance images are added together to get a raw mask $\mathbf{M}_r = \sum_{l=1}^L \mathbf{V}_l^T$.
2. The raw mask image is then thresholded to get a binary mask image:

$$\mathbf{M}_b(x, y) = \begin{cases} 1 & \text{if } \mathbf{M}_r(x, y) > r_t L \\ 0 & \text{else} \end{cases} \quad (6)$$

The threshold ratio r_t determines in which fraction of the L single variance images a pixel has to be present to be considered in the mask. It was fixed at $r_t = 0.25$ in our implementation. It is noteworthy that this value is rather uncritical for the results: in some experiments, $0.15 \leq r_t \leq 0.5$ was found to be valid. Note also that this two-stage thresholding procedure is necessary for success. It is not feasible to just add the variance images \mathbf{V}_l as their dynamic range may be quite different.

3. The binary mask \mathbf{M}_b now approximately outlines the maximum projection of the heart in the image stack but still contains noise clusters as well. However the heart corresponds to the largest cluster in the image. All pixels not corresponding to the largest object are thus set to zero, resulting in the final binary mask \mathbf{M} .
4. The cleaning operator (5) is now applied to \mathbf{M} in order to slightly dilate the outer contour and fill in small holes[‡]. This cleaning step is not strictly necessary and could be left out but was found beneficial for the robustness of the algorithm.

The generation of the mask \mathbf{M} is illustrated in Figure 5. By point-wise multiplication with \mathbf{M} , the heart region is extracted out of the thresholded variance images \mathbf{V}_l^T . The resulting binary images[§] \mathbf{V}_l^M are shown in Figure 6 for the same spatial slice locations as in Figure 3. They are then used to estimate the position, size, aspect ratio and orientation of the heart and match an initial contour by affine transform as described in Section 3. Eventually some more morphological operations to close holes and remove small objects could be performed. However experimentations showed that this is generally not necessary.



Figure 5: Binary mask for the heart region. From left to right: summation of binary variance images, thresholded raw mask, and largest object selected as final mask



Figure 6: Masked variance images for different spatial slice locations

[‡]The parameters used here are 5×5 neighborhood, $n_{\max} = 10$ and $n_{\min} = 7$.

[§]Note that here, the superscript notation T stands for “thresholded” variance image, while M stands for “thresholded and masked” variance images containing only the heart. There are L different \mathbf{V}_l^M .

2.3 Adapting the algorithm to “2D + time” datasets

The algorithm described in Sections 2.1 and 2.2 does not strictly rely on a full 4D dataset with $L > 1$. In the case of $L = 1$, the binary mask before selecting the largest object is identical to the thresholded variance image \mathbf{V}^T of the single spatial slice location and the object corresponding to the heart can be directly selected. However the selected object must really be the heart. This is not necessarily guaranteed in the case of “2D + time” data[¶], so that additional information has to be used. A very simplistic approach that nevertheless has been found to be quite effective is to multiply the variance image \mathbf{V} with a real valued *a priori* confidence mask \mathbf{M}_p before thresholding as in the preceding case. A two dimensional Gaussian function

$$\mathbf{M}_p = e^{-\frac{1}{2\sigma_H}(\mathbf{x}-\mu_p)^T(\mathbf{x}-\mu_p)}$$

is suitable. For the *a priori* mean μ_p the middle of the image $[X/2; Y/2]^T$ can be used in absence of other knowledge. Doing so, we successfully detected the heart in all available “2D + time” image sets.

3 EFFICIENT INITIALIZATION

The binary images \mathbf{V}_i^M obtained in Section 2.2 or 2.3 are next used to adapt a contour model of the heart by affine transform, resulting in valid initial conditions for fully automatic fine segmentation by active contours, further described in Section 4. The following implementation uses 2D discrete-point contour templates, defined as *initial contours* in Section 3.1, which are independently adapted to each slice as described in Section 3.2. This has the advantage of low complexity and is suited for sparse datasets with large inter slice distances^{||}. An extension to 3D models, as used by Park¹⁰ or Terzopoulos⁹ is straightforward however.

3.1 Initial contours

An initial contour^{**} \mathbf{C}_i^I can be obtained from various approaches. The following ones were considered, implemented and compared for the purpose of this work:

- Model-based initialization using a generic or specific model of the heart. Two dimensional templates can be obtained by intersecting three dimensional models.
- Using the results of previous segmentations as initial contours, either the already extracted contour of a neighboring slice or a set of contours from a previously acquired MRI dataset with comparable parameters.
- Automatic initialization as unit circle, which is equivalent to modeling the heart cross sections as ellipses. This simple approach has been found to be very accurate and effective. Drawbacks are that there is no control on either inclusion or exclusion of smaller features –such as the aorta for transversal images– and that in some situations –namely *sagittal* and *long-axis* views of the heart– ellipses are not the best model.

3.2 Adaptation of the initial contours

To adapt an initial contour \mathbf{C}_i^I to each of the binary images \mathbf{V}_i^M , the location, size, orientation and aspect ratio of the heart need first to be estimated from the images. The location, expressed by the center of gravity, is

[¶]In fact, in the lowest and topmost images of the dataset, the heart is not necessarily the largest object in thresholded variance images, as can be seen in left images of Figures 3 and 4.

^{||}Large in comparison to the inter pixel distance

^{**}Accordingly to the previously defined notation of \mathbf{V}_i^M , the superscript notation I of \mathbf{C}_i^I stands for “initial contour”. There are L different initial contours: one for each spatial slice location l .

computed as

$$\mu_{\mathbf{V}_i} = \frac{1}{s_{\mathbf{V}_i}} \sum_{x=1}^X \sum_{y=1}^Y \mathbf{x} \mathbf{V}_i^M(x, y) \quad (7)$$

where $s_{\mathbf{V}_i}$ is the estimated size of the heart –defined as the number of non-zero pixels– and $\mathbf{x} = [x, y]^T$. For the estimation of the orientation and aspect ratio of the heart, we first need to compute a 2×2 covariance matrix

$$\Sigma_{\mathbf{V}_i^M} = \frac{1}{s_{\mathbf{V}_i} - 1} \sum_{x=1}^X \sum_{y=1}^Y (\mathbf{x}\mathbf{x}^T - \mu_{\mathbf{V}_i} \mu_{\mathbf{V}_i}^T) \mathbf{V}_i^M(x, y) = \mathbf{M}_{\mathbf{V}_i}^T \Lambda_{\mathbf{V}_i} \mathbf{M}_{\mathbf{V}_i}. \quad (8)$$

The Karhunen-Loeve transformation¹⁵ is then applied to expand the covariance matrix where $\Lambda_{\mathbf{V}_i}$ is the diagonal matrix of the eigenvalues of $\Sigma_{\mathbf{V}_i^M}$ and $\mathbf{M}_{\mathbf{V}_i}$ is the orthonormal matrix of the corresponding eigenvectors. These entities allow to estimate aspect ratio and orientation.

Similar features for size, location, aspect ratio and orientation can be computed for a contour.¹⁶ This leads to matching the initial contours \mathbf{C}_i^I globally to the binary variance images by *affine transform* such that the features of contour and heart are equivalent. This is done separately for each slice position l . Details are omitted here for brevity and can be found in the references.¹⁶

Note that in some cases where detailed model contours can be provided, it can be advantageous not to correct the rotations based on the eigendecomposition of the variances. Instead, *a priori* knowledge about the orientation of the heart, gained from the image file headers, can be used.

The resulting contours \mathbf{C}_i^A –globally adapted to the heart position– are close enough to the heart outline to be a suitable input for an active contour model, as detailed in Section 4. Figure 7 shows some examples for the adaptation of model-templates (top) and automatically generated ellipses (bottom). It is visible that both model-based contours and ellipses match the contours well with only few differences. Notably, the model-contour fits better at the topmost slice. The aorta is visibly excluded by the model-contours –as intended by the authors– whereas the ellipses intersect this vessel near the middle. This predictable exclusion –or inclusion– of smaller features is a further advantage of the use of model-contours.

4 FINE SEGMENTATION USING SNAKES

In this section, a brief focus is put on how active contour models –also called *snakes*– are used for the fine segmentation of the heart contours. It is assumed that the target has been globally localized and contours \mathbf{C}_i^A have been adapted accordingly. These provided contours now need to be locally deformed for final segmentation which is done in an efficient multi-step approach, using the already computed variance and mean images. The intrinsic energy property of the snakes is first resumed in Section 4.1, while the application of this technique is developed in Section 4.2 within a multi-step fine segmentation approach.

4.1 Active contour energy formulation

Snakes, first introduced by Kass *et al.*⁶ as a method for image segmentation and object outlining, are curves with an associated energy functional $E = E_{\text{ext}} + E_{\text{int}}$ typically composed of an *external energy* or image force E_{ext} pulling the curve towards desired features –*e.g.* intensity edges– and an *internal energy* E_{int} expressing smoothness or model constraints. The snake is then driven to the desired outline by (globally or locally) minimizing this energy. Thus energy formulation and initialization are important for performance. Many improvements to the original snakes and their energy have been suggested, often for application to medical imaging problems.^{5,13,4}




Figure 7: Initial contours adapted to the variance images by affine transform. (Top row: model contours; bottom row: initialization as circle. Left to right: three different spatial slice locations)

The external energy of our snake is computed from an edge map generated by applying a Prewitt edge operator¹⁴ to the image after blurring with a Gaussian kernel. The external energy of a contour point is then found as inner product of contour normal and the so-computed image gradient. The internal energy is computed following the affine invariant model-based approach of Lai^{8,7} but extending it by the addition of a smoothing term and inter-contour relations.¹⁶

4.2 Multi-step fine segmentation of the heart contours

To increase segmentation robustness and computational effectiveness, a multi-step approach was preferred for minimizing the snake energy, utilizing different images and snake performance parameters for each step:

1. In a first step, the globally adapted initial contours \mathbf{C}_l^A are locally deformed using the already computed *binary variance images* \mathbf{V}_l^H for external energy computation. This step performs major deformations necessary to adapt the form of the initial contour \mathbf{C}_l^A –which is already globally adapted by the affine transform– to the outline of the individual imaged heart. As adaption to fine details is not needed at this point, emphasis was put on internal energy, and a relatively large search space was chosen accordingly.
2. The temporal *mean images* \mathbf{M}_l are then used in a second minimization step. This step is aimed to closely follow the contour of the heart. The mean images \mathbf{M}_l are used because small details and spurious edges are not present there, enhancing the robustness of the algorithm. The search space is still relatively large, emphasizing the internal energy.
3. The last step exploits the single images \mathbf{I}_l^t –for all spatial locations l and temporal locations t – as external energy resource for the further minimization of the snakes resulting from the preceding steps, tracking the heart wall motions. For optimization of search efficiency, we start with one image \mathbf{I}_l^t and extract contour \mathbf{C}_l^t , using the result of step 2 as initialization. This contour \mathbf{C}_l^t is then used as initialization on \mathbf{I}_l^{t+1} for

extraction of C_i^{t+1} , which is then used as initialization on the next image and so on. As the contours already match closely, the search space can be reduced.

This multi-step approach allows the use of generic initial contours –ellipses– which are then step-by-step adapted to the image features. If model contours are available, the internal energy can be emphasized and the search space reduced. This method is efficient as the first two steps with larger search region can be performed with a fairly small number of contour points which then can be up-sampled before the third step using spline techniques. Furthermore, this contour extraction technique is robust and fully automatic. It also allows an implementation in which the user can interact and correct the position of single points –*e.g.* using a mouse– if desired. This was found particularly useful after the second step for the lowest and topmost spatial slice locations which exhibit much clutter.

Figure 8 shows the resulting contours on the images for a specific spatial slice location after completion of the different steps. The tracking of the final contour done in the third step can also be very well demonstrated by a computer animation displaying as a movie the images with contours projected on them.




Figure 8: Snake minimized on a binary variance image (left), on the corresponding mean image (middle) and on an actually acquired image (right)

5 CONCLUSIONS

This paper presented a novel approach to the localization and the contour extraction of the human cardiovascular system from a 4D MRI dataset. The key-points of this method are its ability to use all the motion information provided by the data for localization purpose, to use model-based templates or simple generic templates for a coarse initialization of the contours, and to efficiently extract accurate heart contours in a multi-step segmentation process.

An interesting feature of this method is its ability to extract the contours of the heart independently of any user interactions. This was up-to-now a pending challenge for physicians, as the existing techniques often require an initialization stage mainly done by the user. The proposed automatic technique copes with this initialization stage, while taking advantage of existing active contour techniques.

Results of heart wall contour extraction were successfully obtained on transverse 4D MRI datasets. Future works will focus on the extension of the whole technique to 3D and 4D active contours, as well as the introduction of new features such as organ tracking and the common development of visualization tools.

6 ACKNOWLEDGMENTS

Both authors would like to express special thanks to Dr. Luc Bidaut –head of the Digital Imaging Unit of the University Hospital of Geneva, Switzerland– for providing the MR images and guiding them through this captivating research.

7 REFERENCES

- [1] D. H. Ballard. “Generalizing the Hough Transform to detect arbitrary shapes”. *Pattern Recognition*, Vol. 13, pp. 111 – 122, 1981.
- [2] E.M. Brummer, R.J. Lewine, et al. “Automatic detection of brain contours in MRI data sets”. *IEEE Transactions on Medical Imaging*, Vol. 12, No. 2, pp. 153 – 166, June 1993.
- [3] D. Chien and R. R. Edelman. “Fast magnetic resonance imaging”. In Ch. B. Higgins et al., editors, *Magnetic Resonance Imaging of the Body*, Chapter 10, pp. 175 – 198. Raven Press, 2. edition, 1992.
- [4] G. I. Chiou and J.-N. Hwang. “A neural network-based stochastic active contour model (nss-snake) for contour finding of distinct features”. *IEEE Transactions on Image Processing*, Vol. 4, No. 10, pp. 1407 – 1416, October 1995.
- [5] L. Cohen. “On active contour models and balloons”. *CVGIP: Image Understanding*, Vol. 53, No. 2, pp. 211 – 218, March 1991.
- [6] M. Kass, A. Witkin, and D. Terzopoulos. “Snakes: Active contour models”. In *1987 IEEE First International Conference on Computer Vision*, pp. 259 – 268, June 1987.
- [7] K. F. Lai. *Deformable Contours: Modeling, Extraction, Detection and Classification*. PhD thesis, University of Wisconsin-Madison, 1994.
- [8] K. F. Lai and R. T. Chin. “Deformable contours: Modelling and extraction”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 17, No. 11, pp. 1084 – 1090, November 1995.
- [9] T. McInerney and D. Terzopoulos. “A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis”. *Computerized Medical Imaging and Graphics, Special Issue on Cardiopulmonary Imaging*, Vol. 19, No. 1, pp. 69–83, January 1995.
- [10] J. Park, D. Metaxas, A. A. Young, and L. Axel. “Deformable models with parameter functions for cardiac motion analysis from tagged MRI data”. *IEEE Transactions on Medical Imaging*, Vol. 15, No. 3, pp. 290 – 298, June 1996.
- [11] K.P. Philip, E.L. Dove, D.D. McPherson, N.L. Gotteiner, W. Stanford, and K.B. Chandran. “Automatic detection of myocardial contours in cine-computed tomographic images”. *IEEE Transactions on Medical Imaging*, Vol. 13, No. 2, pp. 241 – 253, June 1994.
- [12] K.P. Philip, E.L. Dove, D.D. McPherson, N.L. Gotteiner, W. Stanford, and K.B. Chandran. “The fuzzy Hough Transform feature extraction in medical images”. *IEEE Transactions on Medical Imaging*, Vol. 13, No. 2, pp. 235 – 240, June 1994.
- [13] P. Radeva, J. Serrat, and E. Marti. “A snake for model-based segmentation”. *ICCV*, , No. 8, pp. 816 – 821, 1995.
- [14] A. Rosenfeld. “Image analysis”. In M. P. Ekstrom, editor, *Digital Image Processing Techniques*, Chapter 7, pp. 257 – 287. Academic Press, 1983.
- [15] R.J. Schalkoff. *Pattern Recognition*. John Wiley & Sons, 1992.
- [16] Wolfgang Sörgel. “Model based segmentation of 4D MR images of the heart”. Master’s thesis, University of Erlangen-Nürnberg, Germany, <http://www.nt.e-technik.uni-erlangen.de/~wsoergel/publications/dipl.ps.gz>, 1996.