Progressively Refined Patient-Specific Vessel System Models from Generic Representations

Rongxin Li, Stephen Brown, Laurence Wilson, Jeanne Young and Suhuai Luo
CSIRO Telecommunication and Industrial Physics, PO Box 76, Epping NSW 1710, Australia.
rli@tip.csiro.au

Abstract

In the context of developing visualisation and measurement tools for surgical planning, we have developed systems that construct progressively refined patient-specific vessel system models from multi-level generic representations. In this paper we present a description of our systems, with emphases on (1) knowledge based vessel identification, (2) vessel tracking for linear modelling and (3) injection of human knowledge into the segmentation process for greater accuracy and reliability.

1. Introduction

1.1. Background

Abdominal aortic aneurysms (AAA) are caused by a weakening of the arterial wall, resulting in dilatation and, possibly, eventual rupture of the vessel (see Figure 1). Treatment of abdominal aortic aneurysms increasingly relies on endoluminal procedures (i.e. minimal-access insertion of stent graft) planned using contrast-enhanced X-Ray Computed Tomography (CT) scans [1]. In the planning stage surgeons often need to mentally evaluate multiple CT slices in order to formulate a three-dimensional (3D) representation of the vessel, which must be taken into account when obtaining critical measurements and planning the surgery. Such a practice requires considerable time and experience, and cannot lead to quantitatively accurate results. Due to tortuosity of the vessels, precise measurements of their shape and dimensions are possible only through the use of 3D methods [2], but conversion of the set of CT slices into a three-dimensional model usually involves a large amount of manual intervention. To assist clinicians planning AAA surgery, we have been developing automatic modelling systems (one is fully automatic and the other incorporates user-driven components). The approach we take is a progressive construction of a patient-specific model based on multi-level generic models.

1.2. Overview

In the first stage, the vessel system is modelled on a coarse basis to produce skeleton (or linear) models of vessels that contain only the identity, connectivity and location (and local orientation) information. All suitable curvi-linear entities in the CT data are reduced to 3D linear representations (called skeleton segments) and become candidates for the vessel objects depicted in Figure 1, and their properties (e.g. sizes) and the spatial relationships among them are determined. A knowledge-based inferencing system then compares those properties and relationships with the constraints in a generic model, in order to identify those skeleton segments of interest. The identified skeleton segments are the dimensionality-reduced models of the corresponding vessels. The underlying framework for the knowledge-based inferencing has been applied to two other applications and has been reported in [3], although the current application at that time was largely incomplete and was only briefly described.

Alternatively, the first-stage modelling can be performed with limited help from a human user, in which case the knowledge-based inferencing system will not be involved. The first step is key point matching, in which key points (e.g. bifurcations) in the model are identified in the data by the human user. Next, vessels are linearly modelled by tracking bright curvi-linear structures from the key points.

In the second stage, details of individual vessels are modeled using a dynamic 3D model – with *a priori* built-in generic vessel characteristics – that seeks features in the CT data to evolve into a patient specific model. Part of the processing in this stage has been reported previously [4]. Since then there have been some significant improvements.

The final, optional stage is refinement, where more accurate and reliable results can be achieved as a result of the user's knowledge of the vessel shapes being incorporated as a guiding force of the model's deformation.

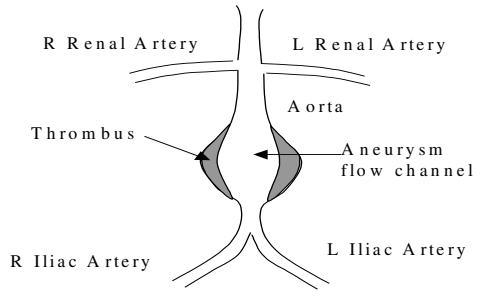

Figure 1 Generic model of the AAA and related arteries

2. Skeleton representation of the vessel system

2.1. Knowledge-based identification

We have designed a knowledge-based model-driven system to identify anatomical structures. The underlying model consists of generic information about the structures to be identified in the image. While the model drives the segmentation, it is expressed separately from the image processing algorithms. The separation of the representation of knowledge from image analysis procedures is the most critical aspect that distinguishes this part of our system from many other image analysis systems that incorporate "knowledge".

Due to the use of contrast medium during the CT imaging, a thresholding operation can separate the arterial network and bone from soft tissue. The inferencing system is used to distinguish contrast-enhanced arteries from bony structures, and to identify (label) arterial segments of interest, notably aorta, the segment subject to aneurysm, iliac arteries and renal arteries. The identification is done by creating a set of correspondences between model objects (i.e. organ labels) and candidates by choosing the "best" candidate for each object in the model.

First, linear representations of curvilinear structures, which include both the vessels of interest and some irrelevant entities, are obtained by a skeletonising operation, consisting of successive erosions of the voxels. These representations are referred to as "skeletons" of the curvilinear structures. Candidates to be matched to vessels in the generic model (see Figure 1) are further extracted from these representations by breaking the skeletonised volume at each branch point. Up to 5 connecting candidates can be merged to cope with the possibility that the aorta may be broken into multiple candidates due to small vessels branching from the aorta, as well as imperfections in the skeletonisation algorithm.

Both unary and binary properties are used in the matching process. Unary properties are intrinsic ones of the candidates. The unary constraints used are the length, orientation and radius of the candidate segments. The length and orientation are calculated directly from the skeletons. The radius is calculated as part of the skeletonisation processing. Binary properties describe the relationship between two candidates, such as relative positions, intensities and sizes. The only binary constraint used here is connectivity between the end points of the aorta, iliacs and renal arteries, with clear distinction between connections at *proximal* and at *distal* ends of vessel segments.

For a given unary or binary property r we define a trapezoidal function $Z(r)$ which defines the "goodness of fit" of the value r to a model value. $Z(r)$ represents the fuzzy membership of the candidate in the set of suitable candidates for a model element. Using the laws of fuzzy logic, the confidence associated with matching an image primitive to a model element M is given by the minimum goodness of fit indicated by the primitive's properties. The best candidate for each structure in the current model subset is then identified by maximising the confidence score.

2.2. Tracking vessels

Whereas identifying large vessels is usually easy for the human expert, the variability of the human anatomy and possible abnormalities pose a challenge to an automated reasoning system. While success can be achieved on many cases, it is difficult to predict the same degree of success on future cases, as some parameters in the system have to be heuristically chosen. Moreover, sometimes the strength of computer assisted surgical planning is not that it can help human experts identify macro anatomical structures, but predominantly lies in the fact that a computerised system can fill in the details between the structures with great speed and accuracy, and produce a visible 3D model.

Therefore, we provide an alternative method for the initial vessel identification, where the human user first identifies a few key points, such as the top of the abdominal aorta and the aorta bifurcation point. After this initial input, the automatic system takes over. The system first expands the identification by vessel tracking.

Our method for tracking vessels is closely related to algorithms used for determining central vessel axes. Those algorithms have been classified into two categories. The first category, indirect approaches, is characterised by the need to pre-process the vessel. The pre-processing can be segmentation of the vessel, in which case the processing flow is the reverse of ours. A more attractive approach is the model-based scale-space filtering, where Gaussian cylinders are used as a model for curvi-linear vessels. After a pre-selection based on the eigenvalues of their local Hessian matrices, local maxima in the scale space of a derivative response function are extracted, which, after further simplification and linking, form vessels' medial axes.

One of the very few direct approaches is proposed by Wink et al [19]. In that approach, the user must define two starting points on the vessel's medial axis. The next point is then iteratively extrapolated and adjusted. In the adjustment, a plane perpendicular to the newly extrapolated medial axis is constructed, wherein the centre of the vessel is determined by searching a region around the extrapolated point for the maximum centre likelihood. The centre is then used as a new point on the medial axis. Drawbacks of this approach include that the tracking process can unpredictably make a wrong turn if the size of the vessel changes quickly.

In some aspects similar to the method proposed by Wink et al, our scheme is also a direct approach. First, a vessel-following "ball" is constructed at the position identified by the user. The ball should be larger than twice the expected

size of the vessel, and should have an initial preferred direction of tracking. From the centre of the ball, scouting elements travel in all directions within a range (e.g. 90°) of the preferred direction, until they are no longer on the foreground of the thresholded data. The orientations associated with the longest distances travelled determine the ball's next position (and the position of the next vessel skeleton point) and the new preferred tracking direction, by an averaging process or by an "inertia-based" method. In the "inertia-based" method, the orientation chosen is the one closest to the latest preferred tracking direction. While the inertia-based method seems to be more robust where the vessel branches, measures must be taken to eliminate the possibility that small errors may be magnified progressively. The measures that we use include error prevention by immediately adjusting the new position of the ball (and that of the tracked point) so that it is at the centre of a foreground component. Another measure is error detection, where the system responds immediately to the situation where the maximum distance travelled by the detecting elements is too small. The system's response is to attempt to re-track the two most recently computed positions using the alternative method.

3. Modelling the Vessels

In indirect approaches, models of vessels are built after they are segmented first. As summarised in [5], traditional segmentation techniques are based on either edge or region information. Edge-based segmentation methods delineate regions by locating high gradient points as edges. These methods are sensitive to noise, and encounter difficulties between regions with small contrast differences. Often edge-linking is a necessary post-processing procedure. Region-based methods, on the other hand, group together homogeneous parts of an image according to certain uniformity criteria. The most widely used method in this category is perhaps region growing. Such methods also often need post-processing using domain knowledge (e.g. [6] to deal with under- and over-segmentation.

A robust segmentation system must incorporate a model-driven aspect, and deformable model approaches are effective ways of doing this [14]. One of the most widely used type of deformable models is active contours (or "snakes") [7] and active surfaces [16][17]. An active surface is a deformable model that evolves in a 3D image under forces that both impose regularity of the surface and attract the surface towards desirable features in the image (eg. gray-level discontinuities). Segmentation using an active surface is therefore both model-driven and data-driven. The outcome, when the forces arrive at the equilibrium, is a surface that is congruent with nearby image features and satisfies some requirements about the surface properties (eg. smoothness). One of the advantages of an active surface model is that it can seamlessly integrate characteristics of the surface with evidence in the image where available.

The model-driven aspect makes the segmentation system more robust. In the case of segmenting the AAA, for example, the model-driven aspect helps overcome potential under-segmentation for the flow channel and the outer surface of the thrombosed region. Under-segmentation is a potential problem when the CT number (called Hounsfield number) for thrombus is similar to that of surrounding soft tissue[1], or when the flow channel is in contact with the spine. In addition to being robust, an active surface model is particularly suitable for the outer surface of an aortic aneurysm since the high internal pressure naturally produces a low curvature surface.

The family of deformable models now also includes a few newer members. The active shape model (and the active appearance model) of Cootes [20] are a notable example. Tek and Kimia [21] have introduced a unified approach to "balloon" techniques using the concepts of "shock fronts" and level sets, linking balloons closely to medial axes, and providing a framework in which reaction-diffusion based on image features plays a more significant role than the physical properties of the balloons. The use of level sets or similar techniques [22][23] can sometimes allow sub-spaces to evolve through topological discontinuities (although not in a way guided by anatomical knowledge).

We use an active-surface vessel model to both segment a vessel and model it at the same time [4]. That is, the model both segments and models the vessel by locating its boundary. This is in fact an important difference from the more traditional segmentation methods.

3.1. Basic model

We have chosen what has been referred to as the constrained 3D model [15] [17]. It is also structurally similar to the model used in [9]. To model the characteristics of vessels, the tubular structure is built into the active surface as a hard constraint.

In a space of permissible deformations, we attempt to minimize an 'energy' functional of the parametric active surface **v** [4] (if desirable the first and second terms in the formula given in [4] may be removed to avoid dependence on the contour parameterization and achieve better invariance to magnification). A minimum of the functional can be obtained by applying the principles of Lagrangian mechanics to construct a corresponding dynamic system [14]. After the system is discretized in both time and space [17], a numerical solution is obtained by iteratively solving and updating a group of linear equations for **v.** The linear equations, which we will refer to in Section 4, can be denoted in the matrix form as

$$\mathbf{A}\,\mathbf{v} = \mathbf{b} \qquad (1)$$

Note that while **A** is a constant matrix determined by the

[1] After formation of the aneurysm, thrombus forms in the outer parts of the dilated segment, preventing contrast medium from entering these regions.

model parameters, both **v** and **b** must be updated in each iteration.

Initially the active surface is merely a generic vessel model. That is, only the generic characteristics that a vessel is a smooth, tubular structure are built into the model. After skeleton vessel models become available as described in the previous section, the active surface then incorporates more specific details of the vessel, namely, the locations and local orientations along the vessel. The active surface is now a tube wrapped around the skeleton vessel model. A complete patient-specific model of the vessel is achieved at the end of the evolution described above.

3.2. Preprocessing of Data

The data provide the external driving force for the model. Unlike those applications that directly convert the image to the external energy in a single step, we employ a re-sampling stage. The CT data is re-sampled at each tracked (for the interactive system) or "skeleton" (for the fully automatic system) point to construct a slice perpendicular to the centerline at that point. These slices are then assembled to form a new data cube, wherein an active surface evolves.

The re-sampling enables the constrained tubular model to be easily used. The corresponding evolution is thus vastly simpler than that inside the original data or that of models that ignore the knowledge of the tubular structure. An added inconvenience, however, is that the resultant model corresponds to the transformed data, and thus a transformation of the model is needed at the end of the evolution.

We have previously shown that it is desirable that the incoming CT data be thresholded [4]. In our experience the threshold does not need to vary from one data set to another.

4. Adding a User-Driven Aspect

The idea of minimal manual intervention after automatic segmentation has been investigated with a neural network based approach [10].

Regarding the active contour/surface approach, there are two basic problems in applying the original formulation. One is that the energy functional is not convex and is therefore vulnerable to local minima [4] [11]. A second problem is the lack of flexibility and adaptability of the model in terms of its topology and regularity.

Since the introduction of the classic active contour model [7], both the data-driven and the model-driven aspects have been tackled in attempts to overcome the above problems. To improve the data-driven aspect of the models, for instance, feature-based external energies have been used, which can be more selective than the original, image-based energy (e.g. [8][12]). Pardo et al [5] have used an elaborate scheme, in which the external energy is composed of multiple parts, incorporating both edge-based and region based terms. Further, irrelevant edge-based information is suppressed through a "focusing" stage, where the strength of the edge and positions of edges in a nearby slice are taken into account.

The model-driven component can be improved by imposing different regularisation constraints (e.g. [13]) or by incorporating extra forces, such as one that initially "inflates" the surface or contour [16] in a bid to overcome some initial local minima. This inflation force is also adopted in our system.

We also developed a few techniques to alleviate the first problem [4]. For the second problem, the prior processing in the first stage largely ensures that the topology of each vessel segment will not change.

In this paper we propose that incorporating another driving force, provided by the user, is a way for further improvement. That is, more accurate and reliable results can result from utilising the user's knowledge of the vessel's shape as a guiding force of the model's evolution. Together with user initiated tracking, the addition of a user-driven aspect enables the system to have a high degree of interactivity. The interactivity in this context has two aspects: the amount of control over the system and its outcome, and the speed with which the system provides feedback. Our interactive system allows the user to control any proportion of the outcome and to have almost instant feedback.

The gray-level discontinuity, crucial to a deformable model, can sometimes be minimal or even non-existent for an entity in a CT image. This is especially true for entities that are not artificially contrast-enhanced, such as the thrombosed region. In such a case, strong model-driven components give rise to smooth segmentation that in effect is not unlike a subjective contour. Unfortunately, an inevitable side effect is that occasionally legitimate image features are treated in the same way as noise, and are ignored.

For example, where an entity has extremely high curvature points, as illustrated in the figure below, some form of intervention or direction by the user may be desired to ensure the accuracy of the fitted surface.

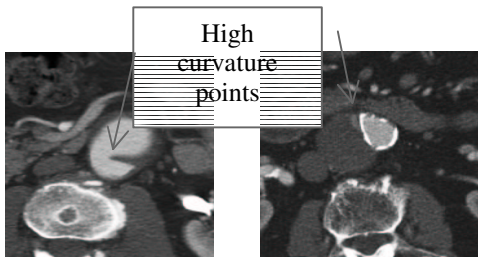

Figure 2 High curvature points on flow-channel surfaces

One scenario of the user providing a driving force to the segmentation process may be that the user selectively marks a few points that the resultant surface must go through. A natural and straightforward way to implement

the driving force in such a case would be to mutate the control point vector **v,** as well as the matrix **A** and the vector **b** in Eq. 1, prior to solving the linear equations. Suppose $\{v_{i_k}, k = 1, 2, \ldots, m , i_k \in [1, N]\}$ are selected to correspond to the marked points, where N is the total number of control points, and m the number of the marked points. The effects of those points upon other points then become *a priori* knowledge and can thus be merged into the right hand side. Then, $\{v_{i_k}, k = 1, 2, \ldots, m , i_k \in [1, N]\}$, along with $\{A_{i_k,j}, k = 1, 2, \ldots, m, i_k \in [1, N], j = 1, 2, \ldots, N\}$, $\{A_{i,i_k}, i = 1, 2, \ldots, N, k = 1, 2, \ldots, m, i_k \in [1, N]\}$[2] and $\{ b_{i_k}, k = 1, 2, \ldots, m\}$ are taken off the equation. Side effects of these mutations are that the dynamic system is irreversibly modified, and that the system no longer corresponds to the surface except through the history of the mutations, both of which may affect the integrity of the system and future operations.

We use an alternative method that avoids the above disadvantages. After the effects of the marked points upon other points are shifted to the right hand side of the equation, **A** is modified such that

$$\left\{ A_{i_k,j} = \begin{cases} b_{i_k}/v_{i_k} & \text{iff } j = i_k \\ 0 & \text{otherwise} \end{cases} , k = 1, 2, \ldots, m, i_k \in [1, N], j = 1, 2, \ldots, N \right\}, \text{ and}$$

$$\left\{ A_{i,i_k} = \begin{cases} b_{i_k}/v_{i_k} & \text{iff } i = i_k \\ 0 & \text{otherwise} \end{cases} , i = 1, 2, \ldots, N, k = 1, 2, \ldots, m, i_k \in [1, N] \right\}$$

Although b_{i_k} changes in each iteration, this method is not inefficient as in practice it is not necessary to calculate b_{i_k}/v_{i_k}.

5. Implementation and Preliminary Evaluation

At present the knowledge-based and the interactive approaches for the first modelling stage are implemented in two separate systems, although they will be integrated in a future version. The interactive system is not yet complete. Our systems are implemented using Java™, which provides not only platform independence and convenient visualisation of 3D patient-specific models, but also the possibility to easily convert the system to a server-client suite for network-centric applications. Additional functions, such as observing the surface deforming, are implemented for education and demonstration purposes.

Our system has not yet advanced to the formal evaluation stage. Consequently, we can only present some preliminary evaluation results here. In Figure 3, the green and pink coloured surfaces (for the flow channel and aorta outer surface respectively) were output by the fully automatic version of our system, whereas the red and yellow coloured ones were produced by experienced radiographers. Each pair of surfaces are so close that either one is alternately obscuring the other. Figure 4 shows a screen shot of the interactive version of our system. On the top left, slider bars provide mechanisms for easy adjustment of the model parameters. Below it is the main window, where the user directly interacts with the data and the model. The vessel models can be examined slice by slice in the main window. The screenshot includes such a slice, with projections of the model surfaces shown as red contours. The user can select interaction modes from the menu or by right-clicking the mouse inside the window. On the right, a 3D patient-specific model and various critical measurements calculated based on the model are displayed (the measurements and algorithms used for calculating them are beyond the scope of this paper). In Figure 5 a few examples segmented by the fully automatic version of our system are displayed.

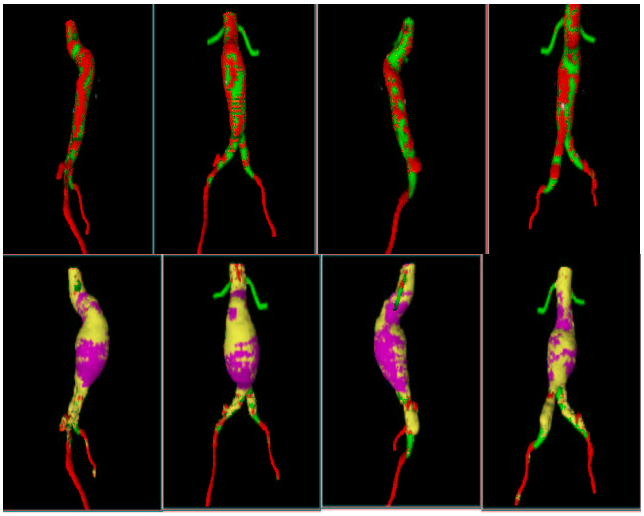

Figure 3 Comparison with manual segmentation

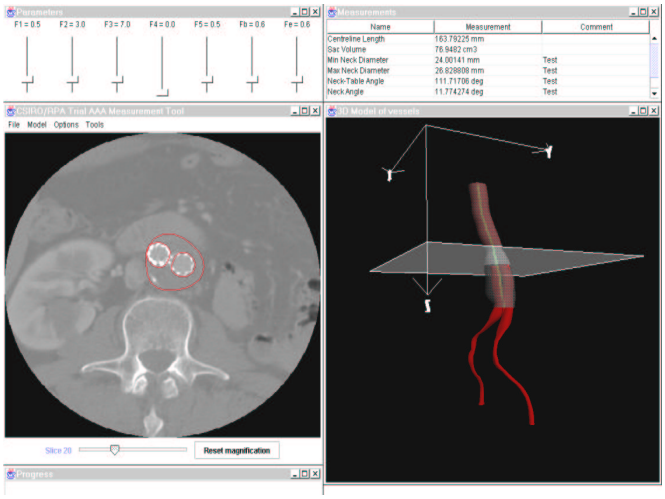

Figure 4 Screenshot of the interactive system

[2] Note that the two sets are not disjoint.

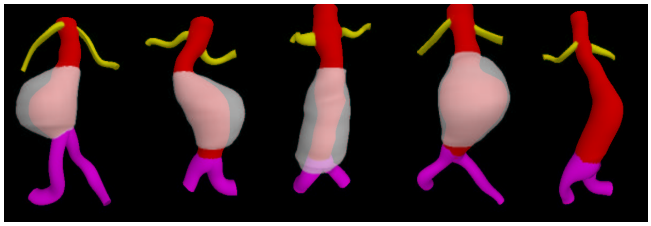

Figure 5. Some results produced by the fully automatic version of our system

6. Conclusion

In the context of developing visualization and measurement tools for surgical planning, we have developed systems that construct progressively finer patient-specific vessel-system models from multi-level generic representations. First, knowledge-based vessel-matching or human-initiated vessel-tracking converts the generic relational model to a patient-specific vessel skeleton model, which further incorporates location and orientation information. Next, details of the vessel surfaces are further modeled with a dynamic vessel model. Finally, a user-driven aspect can be optionally added to the evolution of the vessel model for greater reliability or refinement. We have shown that such a progressive paradigm is highly effective for the segmentation and modeling of a major vessel system.

7. Acknowledgments

The authors thank Dr. Hence Verhagen and Dr. Geoffrey White of Royal Prince Alfred Hospital for their enthusiastic support and cooperation.

The authors are grateful towards Medical Media Systems, Inc. for their kind permission to use the evaluation data.

8. References

[1] May, J., White, G. H., Yu, W., Waugh, R. C., Stephen, M. S., McGahan, T. J., and Harris, J., "Early experience with the Sydney and EVT prostheses for endoluminal treatment of abdominal aortic aneurysms," J. Endovascular Surgery, 2, 240-247, 1995

[2] Fillinger, M. F. New imaging techniques in endovascular surgery. Endovascular and Minimal Invasive Vascular Surgery, 79(3), 451 – 475, 1999.

[3] Wilson, L. S., Brown, S. F., Brown, M. S., Young, J. A., Li, R., Luo, S. and Brandt, L. Segmentation of medical images using explicit anatomical knowledge. In Pham, B., Braun, M., Maeder, A.J. and Eckert, M.P. (Eds.) New Approaches in Medical Image Analysis, Proceedings of SPIE Vol. 3747, 142-155, 1999.

[4] Li, R., Brown, S. F., Wilson, L. S. and Luo, S. Extraction and Reconstruction of Vessel Walls from CT Scans Using a Three-Dimensional Deformable Model. Proc. Digital Image Computing: Techniques and Applications, 138-143, 1999.

[5] X.M. Pardo, M.J. Carreira, A. Mosquera and D. Cabello. A snake for CT image segmentation integrating region and edge information. Image and Vision Computing 19, 461-475, 2001

[6] M. Sonka, W. Park, E.A. Hoffman, Rule-based detection of intrathoracic airway trees, IEEE Transactions on Medical Imaging 1 (3), 314-326, 1996

[7] Kass, M., Witkin, A. and Terzopoulos, D. Snakes: Active contour models, Internat. J. Computer Vision, 1, 259 - 268, 1987

[8] Cohen, L. On active contour models and balloons, CVGIP: Image Understanding 53, 211-218, 1991.

[9] R. Shekhar, R.M. Cothren , D.G. Vince, S. Chandra , J.D. Thomas, J.F. Cornhill. Three-dimensional segmentation of luminal and adventitial borders in serial intravascular ultrasound images. Computerized Medical Imaging and Graphics 23 299–309 1999.

[10] Shiffman, S, Rubin, G. D. and Napel, S. Semiautomated editing of computed tomography sections for visualization of vasculature. , Proceedings of SPIE - The International Society for Optical Engineering, 140-151, 2707,. 1996

[11] Jain, A., Zhong, Y and Dubuisson-Jolly, M-P. Deformable template models: A review. Signal Proc. 71(2), 109-129, 1998.

[12] Abrantes A. and Marques J., A class of constrained clustering algorithms for object boundary extraction, IEEE Trans. Image Process. 5, 1507 - 1521, 1996.

[13] Christensen, G.E., Rabbitt, R.D. and Miller, M.I. Deformable templates using large deformation kinematics, IEEE Trans. Image Process. 5, 1435 - 1447, 1996.

[14] McInerney, T. and Terzopoulos, D. Deformable models in medical image analysis: A survey. Medical Image Analysis, Vol. 1, No. 2, 91 – 108, 1996.

[15] Terzopoulos, D. , Witkin, A. and Kass, M. Constraints on deformable models: Recovering 3D shape and nonrigid motion, Artificial Intelligence, 36(1), 91-123, 1988.

[16] Cohen, I., Cohen, L. D. and Ayache, Y. Using deformable surfaces to segment 3D images and infer differential structures. CVGIP: Image Understanding, 56(2), 242 – 263, 1992.

[17] Cohen, L. D. and Cohen, I. Finite-Element methods for active contour models and balloons for 2D and 3D images. IEEE transactions on Pattern Analysis and Machine Intelligence, 15(11), 1131-1147, 1993.

[18] Bamford, P. and Lovell, B. Unsupervised cell nucleus segmentation with active contours. Signal Processing Special Issue: Deformable Models and Techniques for Image and Signal Processing, 71(2), 203-213, 1998.

[19] Wink, O., Niessen, W. J., and Viergever, M. A., Fast Delineation and Visualization of Vessels in 3-D Angiographic Images, IEEE Trans. Medical Imaging, 19. 337-346, 2000.

[20] Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J., Active shape models, their training and application, Computer Vision and Image Understanding, 61, 38-59, 1995.

[21] Tek, H. and Kimia, B. B., Volumetric segmentation of medical images by three-dimensional bubbles, Computer Vision and Image Understanding,. 65. 246-258, 1997.

[22] Lorigo, L., Faugera, O., Grimson, W., Keriven, R., Kikinis, R., and Westin, C.-F., "Co-dimension 2 geodesic active contours for MRA segmentation," in Kuba, A., Samal, M., and Todd-Pokropek, A. (eds.) *Information Processing in Medical Imaging* Berlin: Springer-Verlag, 1999, 126-139.

[23] McInerney, T. and Terzopoulos, D., Topology adaptive deformable surfaces for medical image volume segmentation, IEEE Transactions on Medical Imaging,. 18(10) 840-850, 1999.