Research Article

Force Feedback to Assist Active Contour Modelling for Tracheal Stenosis Segmentation

Lode Vanacken,[1] Rômulo Pinho,[2] Jan Sijbers,[2] and Karin Coninx[1]

[1] *Hasselt University-tUL-IBBT Expertise Centre for Digital Media, Wetenschapspark 2, 3590 Diepenbeek, Belgium*
[2] *IBBT Vision Lab, Department of Physics, University of Antwerp, Universiteitsplein 1, N.1, 2610 Wilrijk, Belgium*

Correspondence should be addressed to Lode Vanacken, lode.vanacken@uhasselt.be

Manual segmentation of structures for diagnosis and treatment of various diseases is a very time-consuming procedure. Therefore, some level of automation during the segmentation is desired, as it often significantly reduces the segmentation time. A typical solution is to allow manual interaction to steer the segmentation process, which is known as semiautomatic segmentation. In 2D, such interaction is usually achieved with click-and-drag operations, but in 3D a more sophisticated interface is called for. In this paper, we propose a semi-automatic Active Contour Modelling for the delineation of medical structures in 3D, tomographic images. Interaction is implemented with the employment of a 3D haptic device, which is used to steer the contour deformation towards the correct boundaries. In this way, valuable haptic feedback is provided about the 3D surface and its deformation. Experiments on simulated and real tracheal CT data showed that the proposed technique is an intuitive and effective segmentation mechanism.

1. Introduction

Image segmentation in the medical field is an important step for the diagnosis and treatment of various diseases. In many cases, this task is performed manually [1, 2]. However, manual segmentation is widely acknowledged as being time consuming and intra- and interoperator dependent. Hence, some level of automation during the segmentation is desired, as it often significantly reduces the segmentation time. Medical image segmentation, in particular, is a very complex task, given the necessary precision required for object extraction and boundary delineation. A typical solution is to allow users to provide extra knowledge to or interfere with the segmentation process in order to refine the results yielded by the automatic steps, which is known as semiautomatic (or interactive) segmentation.

The Active Contour Model (ACM) [3] is a well-known shape deformation algorithm to delineate structures in images, and several semiautomatic versions of this algorithm have been proposed in the literature [4]. ACMs minimise an energy function that controls the bending and stretching of a given initial contour and the attraction by image features.

The expected result is that the contour matches the boundary of the structure of interest in the image. In 2D, the interface between user and algorithm is usually established with click/drag processes. However, if the data being segmented is three-dimensional, such as in 3D Computed Tomography (CT) images, a more refined interface is called for.

The present work sets forth a 3D segmentation interface for ACMs based on haptics. The chosen application is the segmentation of tracheal stenosis from chest CT scans. Tracheal stenosis is an unnatural narrowing of the trachea with traumatic, neoplastic, or idiopathic causes that, despite being relatively rare, can be life threatening [5]. In order to correctly diagnose and treat it, accurate assessment of the stricture is necessary, which determines the point where it starts, where it ends, and the degree of narrowing [6]. One way of performing this assessment is via segmentation [7], which needs to be especially accurate around the narrowed parts of the trachea, so that the parameters of the stenosis can be correctly calculated. Although the healthy trachea can in general be segmented very easily with, for example, region growing, the task may be challenging in cases of severe stenosis. As discussed in [7], this happens because

the tracheal lumen at the narrowed region is often barely visible in the image. Triglia et al. [8] solved this problem by manually reconstructing the parts of the narrowed trachea that were not visible in the image, but this can be time consuming and prone to error. In [9], it was also shown that 3D ACMs are able to reconstruct missing parts of the trachea, but noise and the presence of neighbour organs in the image may hinder segmentation. The semiautomatic process proposed here is therefore meant to overcome such difficulties.

The proposed ACM can be steered by the user with 3D input from a haptic device. Conversely, the user is provided with valuable force feedback about the 3D surface and its deformation. This interaction creates a first-person 3D environment, which gives the user the feeling that a real shape is being manipulated. The net effect is an intuitive, interactive segmentation mechanism that improves over traditional 2D approaches. The method was evaluated with two sets of 3D CT images. The first is a real case of severe tracheal stenosis. The second set is a phantom of stenosis created from a real CT image, in which the oesophagus also appears in the image and may disturb the segmentation process. The results obtained with the proposed method were compared to a reference manual segmentation using traditional region growing.

This paper is organised as follows. Section 2 gives an overview of existing semiautomatic segmentation algorithms, including those employing 3D interfaces. Section 3.2 briefly reviews the ACM used in the segmentation of tracheal stenosis from CT images. In Section 4, the method to integrate haptics with ACM is fully described. The experiments and evaluation of the proposed method are given in Section 5. Section 6 presents a discussion of the results obtained, and the paper is concluded in Section 7.

2. Related Work

The work related to the research presented here can be categorised in primarily two different topics, namely, interactive segmentation and the addition of force feedback to segmentation. Both will be discussed in this section.

2.1. Interactive Segmentation. Given the difficulties in implementing fully automatic segmentation algorithms, interactive segmentation has long been a topic of interest in the literature. Qiu and Yuen [1] recently discussed current trends and the history of semiautomatic segmentations. Olabarriaga and Smeulders [2] presented an interesting survey focussed on interactive segmentation in medical images, discussing practical and subjective aspects of the problem. McGuinness and O'Connor [10] proposed a framework for the evaluation of interactive segmentation and evaluated four existing algorithms. McInerney and Terzopoulos [4] presented a survey of ACMs applied to medical images and described different ways of manipulating the deformable curves with user intervention.

In the field of segmentation of 3D images, Kang et al. [9] proposed a set of editing tools meant to interactively correct inaccuracies of automatic segmentation methods. They compared their tools with traditional slice-by-slice 2D segmentation approaches. Another similar and successful approach has been proposed by Heckel et al. [11], which applied variational interpolation in combination with a set of user-drawn, planar contours that can be arbitrarily oriented in 3D space. A natural evolution of such types of interaction tools is the use of more advanced, or even 3D, interfaces. Bornik et al. [12], for example, used immersive 3D spaces and tablets to manipulate deformable models in segmentation refinement, and Zudilova-Seinstra et al. [13] recently employed glove-based input for the delineation of medical data.

2.2. Force Feedback during Segmentation. As 3D interfaces evolved to haptics, so did the 3D user interaction in segmentation applications. Vidholm et al. [14, 15] used haptics such that seeds for the segmentation could be placed in good spots. Force feedback was applied based on data from MR images so that the user could recognise good seed locations. Similarly, Malmberg et al. [16] augmented a 3D wiring segmentation technique with force feedback. In this type of segmentation technique, the boundaries of the object to be segmented are contoured (wired). Force feedback using volume haptics was provided to better understand the object boundaries and thus enhance the wiring.

More related to the work presented here, Vidholm et al. [17] showed a technique which enables the user to push a deformable model to perform segmentation. Finally, when working with deformable models, the placement and shape of the initial surface to be deformed is also very important. Harders and Székely [18] proposed to extract the centre line of a tubular structure to create forces to guide a user on a path close to this centre line. While moving along the path, the user sets control points which define a spline-based centre line. This centre line, in turn, is used to calculate and generate a cylinder with dynamic width to represent the deformable model.

3. Segmentation of Tracheal Stenosis

In the following, we will briefly review the main concepts behind ACMs and how they are used to accomplish the tracheal stenosis segmentation.

3.1. Active Contour Models. ACMs, commonly known as snakes, are curves defined within an image domain that are able to move under the influence of internal forces derived from the curve itself and of external forces derived from the image data. The internal and external forces are defined in such a way that the curve will register to an object boundary or other desired features within an image. As defined by Kass et al. [3], a snake can be represented in 2D by a curve $\mathbf{v}(s) = (x(s), y(s))$, $s \in [0, 1]$, responding to an energy functional of the form

$$E = \int_0^1 [\kappa E_{\text{int}}(\mathbf{v}(s)) + (1 - \kappa)E_{\text{ext}}(\mathbf{v}(s))]ds, \quad (1)$$

where $\kappa \in [0, 1]$ is a weighting factor.

(a) (b) (c)

FIGURE 1: (a) The initial shape of the ACM (green) with the corresponding landmarks (red dots). (b) A step of the deformation algorithm. (c) The landmarks of the shape being deformed overlaid on an image of tracheal stenosis. Note the landmarks being attracted by the edges of the structures in the image.

The internal energy E_{int} restricts the deformations by preventing the curve from breaking apart and avoiding the appearance of sharp corners. The external energy usually represents the gradient of an image I convolved with a Gaussian function G at scale σ, which causes the curve to be attracted by contours with high image gradients.

The objective is then to minimise (1), making the system a force balance equation of the form

$$\kappa \mathbf{F}_{\text{int}} + (1 - \kappa) \mathbf{F}_{\text{ext}} = \mathbf{0}, \tag{2}$$

where

$$\begin{aligned} \mathbf{F}_{\text{int}} &= -\nabla E_{\text{int}}, \\ \mathbf{F}_{\text{ext}} &= -\nabla E_{\text{ext}}. \end{aligned} \tag{3}$$

The minimisation is solved iteratively. The expected result is a curve that matches the high gradients of the image while being restricted by the internal constraints, according to the assigned weighting factors.

3.2. ACM for Tracheal Stenosis Segmentation. The concepts above naturally extrapolate to 3D and can easily be adapted to a discrete domain. Within the context of the proposed application, a discrete surface is defined as $\mathcal{S} = (\mathcal{X}, \mathcal{T})$, where $\mathcal{X}$ is the set of points, or landmarks, with $\mathbf{x}_i, i = 1, \ldots, n$, a point in this set, and $\mathcal{T}$ is the set of triangles connecting the points of $\mathcal{X}$.

The ACM is initialised with an estimation of the healthy shape of the trachea, obtained with the method proposed in Pinho et al. [7]. In this way, the initial shape tends to be near enough to the boundary of the narrowed trachea in the image. In addition, this shape conveys more intuitive information to the user. This improves on the approach proposed by Vidholm et al. [17], in which the initial surface does not necessarily resemble the target object.

The deformation algorithm iteratively loops through all the points in $\mathcal{X}$, applying the ACM forces locally, until no significant deformation has been made to the surface. Below, the internal and external forces are briefly presented and Figure 1 illustrates some steps of the algorithm. For a detailed description, we refer the reader to [7].

3.2.1. External Force. The external force $\mathbf{F}_{\text{ext}}$ is derived from the image, which is first converted into a distance map I_D indicating the distance from any point to the nearest edge (high image gradient). The gradient of I_D defines how landmarks of the surface lying on a certain point of the image are influenced by $\mathbf{F}_{\text{ext}}$. Therefore, the external force applied to the landmark $\mathbf{x}_i$ of $\mathcal{X}$ is defined as

$$\mathbf{F}_{\text{ext}_i} = -\frac{|\nabla I_D(\mathbf{x}_i)|}{M} \nabla I_D(\mathbf{x}_i), \tag{4}$$

where M is the maximum gradient magnitude in I_D.

3.2.2. Internal Forces. The internal force $\mathbf{F}_{\text{int}}$ controls stretching and bending, in such a way that the surface is continuous (does not break apart) and remains smooth (has no sharp corners). The force tries to keep the landmarks equally spaced and tries to minimise the local Gaussian curvature of the surface. It is given by

$$\mathbf{F}_{\text{int}_i} = \gamma \mathbf{F}_{\text{elast}_i} + (1 - \gamma) \mathbf{F}_{\text{bend}_i}, \tag{5}$$

where γ is a weighting factor.

$\mathbf{F}_{\text{elast}_i}$ is the elastic force applied to $\mathbf{x}_i$ of $\mathcal{X}$, defined as

$$\mathbf{F}_{\text{elast}_i} = D_i \frac{\mathbf{d}_{\text{elast}_i}}{|\mathbf{d}_{\text{elast}_i}|}, \tag{6}$$

where the directional component $\mathbf{d}_{\text{elast}_i}$ moves the landmark towards a central point relative to its neighbours. The scalar component D_i, in turn, is a normalised measure of how much $\mathbf{x}_i$ deviates from this central point.

The bending force $\mathbf{F}_{\text{bend}_i}$ is given by

$$\mathbf{F}_{\text{bend}_i} = K_{G_i} \frac{\mathbf{d}_{\text{bend}_i}}{|\mathbf{d}_{\text{bend}_i}|}, \tag{7}$$

where $\mathbf{d}_{\text{bend}_i}$ is either equal to $\mathbf{d}_{\text{elast}_i}$ or it moves $\mathbf{x}_i$ along its normal if the landmark is not located at the open ends of the surface. In either case, the directional component moves $\mathbf{x}_i$ in such a way that the discrete Gaussian curvature computed at the landmark is minimised. The scalar component K_{G_i} is a normalised measure of how much the curvature at $\mathbf{x}_i$ deviates from zero.

Finally, for each iteration j of the deformation algorithm of the ACM,

$$\mathbf{x}_i^{(j)} = \mathbf{x}_i^{(j-1)} + \kappa \mathbf{F}_{\text{int}_i} + (1 - \kappa) \mathbf{F}_{\text{ext}_i}. \tag{8}$$

4. Haptic Interaction in Active Contour Models

In this section, we will first give some theoretical background on the use of haptics and the benefits it can provide to 3D interaction. Afterwards, we will discuss how we integrated force feedback with ACMs when they are used to segment tracheal stenosis.

4.1. Haptic Interaction. Haptic interaction requires the existence of a device that serves as the interface between the user and an application. Not only does this device enable the user to provide input to the application, but also conveys to the user the reaction of the system to the provided input (force feedback), giving the user a sense of touch [19].

A first step towards adding force feedback to any application is deciding how to perform the haptic rendering. In our case, the trachea is represented using the surface description $\mathcal{S} = (\mathcal{X}, \mathcal{T})$ of Section 3.2. Using the haptic device, the user can interact with the triangular boundary of $\mathcal{S}$, pushing or pulling it. The haptic rendering takes care of interpreting the commands sent by the haptic device, transferring them to the surface, and returning to the user the force feedback given by the surface. This is performed using the algorithm described by Ruspini et al. [20], illustrated in Figure 2. The force feedback in this algorithm is calculated using Hooke's law:

$$\mathbf{F}_{\text{user}} = -k\mathbf{y}. \tag{9}$$

When rendering force feedback, the force applied by the device to the user is calculated using $k\mathbf{y}$. The user actually slightly penetrates the surface and the distance of this penetration, $\mathbf{y}$, together with a constant k is used to calculate this force. The constant k represents the stiffness of the model and indicates how the surface being touched reacts as a function of $\mathbf{y}$. As a result, $\mathbf{F}_{\text{user}}$ turns out to be as well the force applied by the user to the surface. The duality of $\mathbf{F}_{\text{user}}$ is one of the advantages of haptics and is beneficial to both the system and to the user, increasing the sense of first-person interaction in the 3D environment.


$$\mathbf{F}_{\text{user}} = -k\mathbf{y}$$

Figure 2: Conceptual representation of haptic rendering. k determines the stiffness of the surface and $\mathbf{y}$ the penetration distance between the haptic probe (red circle) and its projection on the surface border (green circle).

4.2. Integration with ACMs. As described in Section 3.1, the external force used in this work for the ACM derives from the gradient of a distance map I_D that indicates the distance from any point in the original image to the nearest edge. In our haptic framework, $\mathbf{F}_{\text{user}}$ acts as an additional external force to the ACM, creating the new force

$$\mathbf{F}_{\text{hap}} = \mathbf{F}_{\text{ext}} \pm \mathbf{F}_{\text{user}}, \tag{10}$$

which changes (2) to

$$\kappa \mathbf{F}_{\text{int}} + (1 - \kappa) \mathbf{F}_{\text{hap}} = \mathbf{0}. \tag{11}$$

Note that $\mathbf{F}_{\text{user}}$ can represent two directions and, depending on its sign, can symbolise to the user either a pushing or pulling gesture.

Making $\mathbf{F}_{\text{user}}$ an additional external force allows the user to interact with the segmentation in real time and provide additional input in order to enhance the segmentation results by correcting the ACM forces in cases in which it was not converging to the correct locations. Nonetheless, one pitfall of this added force is that users might strongly deform the tracheal surface, by either pushing or pulling it in the wrong direction. For instance, once the triangle being touched and its landmarks are identified, adding the force to only those landmarks would punch a hole in the surface and would make it very difficult to influence the ACM more globally. This effect is avoided by smoothly propagating the user force to the neighbour landmarks according to a Gaussian function of the distance from the landmark to the point of contact. Although other smoothing functions could be chosen, for example, a 3rd-order B-Spline, the Gaussian function represented a good compromise between physical correctness and computational cost. In addition, as already explained, the internal forces of the ACM tend to preserve the smoothness and continuity of the surface. As a result, with, for example, $\kappa = 0.8$ in (11), the surface will tend to automatically repair itself following any severe deformation.

In contrast, the weight κ can take a lower value while the user is pushing or pulling the surface. With $\kappa = 0.5$, for instance, the user has more control over the deformation, since the internal forces of the ACM will be relaxed. This weight, however, is only applied to those landmarks being directly affected by the user force, which limits the deformations to a restricted area.

Another feature is the possibility to pause the ACM forces such that the user has complete control over the surface. This gives the user extra time to diagnose areas that are not

converging to the desired locations. Pausing the ACM means that the forces are only applied to landmarks being pushed or pulled, that is, $\mathbf{F}_{\text{hap}_i} = \mathbf{F}_{\text{user}_i}$.

We further augmented the typical force equation with an extra transformation which depends on the gradient of the original image and alters the constant k in (9) to

$$k' = k \frac{|\nabla I_D(\mathbf{x}_i)|}{M}, \qquad (12)$$

with $\nabla I_D(\mathbf{x}_i)$ and M is the maximum gradient magnitude in I_D (as defined in (4)).

If the magnitude of the gradient of the image at the point of contact with the surface is high, the surface will seem more stiff to the user. Conversely, if the magnitude is small the surface will seem more flexible. This change in stiffness provides the user with an interesting variation in the force feedback. Remember that the aim of the external force of the ACM is to guide the surface towards high gradients (edges) of the image. At these locations, the user will have more difficulty in pushing or pulling the surface, meaning that the surface may already be near or resting on the correct place. Yet, if the surface is in reality near a wrong edge, the user can still manipulate the segmentation by increasing the force exerted via the haptic device.

Finally, another aspect to be taken into account is the high update rate of 1 kHz necessary for stable realistic haptics [20]. This constraint does not match well with the ACM algorithm being used for segmentation, as one iterative loop of this algorithm is computationally intensive enough to take longer than one millisecond to run. Therefore, it is necessary to decouple the segmentation, graphics rendering, and haptic rendering in separate parts (threads) such that they can all run in parallel (see Figure 3). This guarantees that the user interaction is not hindered. The graphics thread constantly renders on the screen the new shape which is provided by the ACM thread. The ACM thread performs the ACM segmentation and after every iteration it also provides the haptic shape to the haptic thread. This thread on its turn provides the haptic rendering as well as the force from the user to the ACM thread such that it can be used during the segmentation.

At this point, we have covered all aspects involved in the proposed segmentation of tracheal stenosis using haptics. In the next section, the proposed method will be evaluated through a series of experiments.

5. Experiments

We carried out a number of experiments to evaluate whether the addition of haptics allows users to influence the ACM segmentation adequately. The idea was to collect the user impression of the system and to quantitatively evaluate the segmentation results. Since we would also like to judge the importance of the addition of force feedback, the experiments were divided into executions having force feedback switched on and off.

5.1. Data. We used two CT images in this series of experiments that clearly demonstrate the segmentation challenges
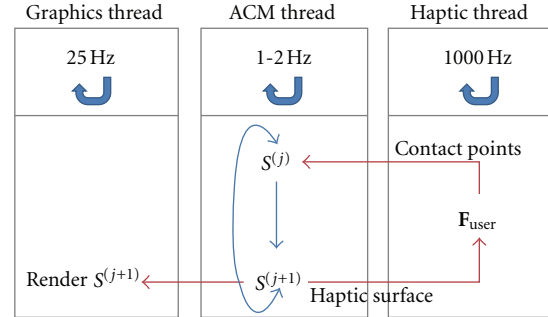

FIGURE 3: Diagram of the separate parts involved in the system. The ACM thread performs the ACM segmentation influenced by the forces provided by the user $\mathbf{F}_{\text{user}_i}$. After performing one iteration j of the segmentation, the ACM thread provides the new shape $\mathcal{S}^{j+1}$ to the graphics rendering and haptic rendering threads.

mentioned in Section 1. The first image was a real case of severe stenosis extending along 2/3 of the trachea. The difficulty in this segmentation, when using traditional ACM, is that, due to the severity of the narrowing, the tracheal lumen at the narrowest location of the trachea is barely visible in the CT image of the patient. As a result, the edges of the tracheal wall are not well defined. With region growing, the segmentation is actually nearly split in two (see Figure 4(a)). Our aim here is to give fine control to the user such that he or she can steer the ACM to the correct locations using the defined forces.

In the second image, a phantom of stenosis was created from the CT scan of a healthy patient. In this scan, both the trachea and the oesophagus, located behind the trachea, are visible. They were first segmented with traditional, semi-automated region growing, generating a binary image. This binary image was further processed with a tool to manually create stenosis in an otherwise healthy trachea. In the tool, the user places an erosion mask on the trachea and iteratively erodes the regions below the mask until the stenosis achieves the desired shape (see Figure 4(b)). The difficulty in this case lies in the fact that the oesophagus is very near the trachea. Although the trachea alone could be segmented with, for example, region growing in this case, one could imagine a situation in which the two tubes would appear connected in the image (due to artefacts or anatomical anomalies), making region growing an inadequate choice. The consequence of such configuration to the ACM is that it directly affects its external force, to such an extent that the contour controlled by the traditional ACM cannot be attracted by the edges of the trachea. The idea is that the user steers the ACM into the capture range of the edges of the trachea, from where the ACM can continue, in principle, with no further interaction.

5.2. Participants. Four volunteers served as participants in this experiment, all of them had at least some experience with virtual environments. Although none of the participants had experience with medical images, one of them was familiar with segmentation algorithms. We judged that familiarity with virtual environments and haptics was more important

(a)   (b)

Figure 4: (a) A real, very severe case of stenosis segmented with simple region growing. (b) The trachea and the oesophagus appearing together in the image, after segmentation, followed by the phantom stenosis created on the same trachea.


Figure 5: The set-up used to perform ACM segmentation for tracheal stenosis using a Phantom premium 1.5.

than with segmentation because our objective was to collect the impression the participants had about the human-computer interaction offered by the system.

5.3. Hardware. For input, two devices were used. As described earlier, we need first and foremost a device to provide force feedback to the users and to allow them to influence the segmentation. A PHANToM premium 1.5 was used for 6 DOF input and 3 DOF force feedback. This device is equipped with a stylus having a single button. The haptic set-up was similar to the one shown in Figure 5.

The second device was used for navigation capabilities and system control. It is required that the shape to be deformed can be viewed and touched from different angles and locations. A space mouse with 6 DOF was thus used for this task. Such devices are typically used by the nondominant hand and are therefore suited to be used in combination with

the PHANToM [21]. Regarding system control, the space mouse's buttons were used to start, pause, restart, finish the segmentation, and to disable or enable the display of visual features.

The computer set-up we used consisted of an Intel Xeon E5520 at 2.27 GHz, with 4 GB RAM and NVIDIA Quaddro FX 4800 graphics card. As display we used a 21-inch-wide screen with a resolution of 1680 by 1050.

5.4. Software. The application developed for the segmentations is responsible for establishing the interaction between the user and segmentation task. It displays to the user the ACM algorithm in action and provides to the algorithm the data received from the input devices. The application starts by displaying the shape to be deformed, already placed at the correct start location. Remember from Section 3.2 that the initial ACM shape is an estimation of the healthy trachea of the patient, obtained with the method proposed in [7]. In this way, the focus of the experiment was not on the placement of the initial shape, but only on the segmentation itself. The landmarks of the ACM shape are also visualised in the application, such that the user has a better view on how the deformations are occurring as the segmentation algorithm iterates.

In order to enhance the users' haptic feedback experience, three visual cues were included in the application. The first is to highlight the triangles being touched by the user. In this way, not only do the users perceive touch through the force feedback, but also they can be sure about the point of interaction. The second visual cue is a line emanating from the interaction point indicating the strength of the force exerted onto the shape. This is especially useful during a training phase, when the users can have a clearer idea about the relationship between the force applied through the haptics device and the actual deformation of the ACM shape. Finally, it is important that the user can see if the segmentation is succeeding in delineating the boundary of the trachea in the CT image. The application therefore integrates the visualisation of the 3D shape of the trachea

FIGURE 6: Visualisation of the interaction point (red) and its projections (yellow, pink and blue) onto the slices as well as the CT texture data on the slices representing the current interaction point position.

TABLE 1: Statistics of segmentation quality.

	Real data			Phantom data		
	Sensitivity	DSC	Jaccard	Sensitivity	DSC	Jaccard
ACM	0.896	0.737	0.583	0.902	0.855	0.747
Hap	0.871	0.685	0.523	0.886	0.850	0.743
Hap (ff on)	0.873	0.684	0.521	0.922	0.874	0.777
Hap (ff off)	0.870	0.687	0.525	0.849	0.826	0.709

with 2D views of the CT image volume. These views are projections of the axial, sagittal, and coronal planes of the image volume. However, instead of the traditional visualisation of preselected slices of the image, the application projects on three fixed planes the slices corresponding to the closest integers derived from the coordinates of the interaction point (see Figure 6). Such feature provides the user with a clear indication of how the image data looks like at the current position of the haptic device.

5.5. Procedure. A repeated measures within-participant design was chosen to evaluate the proposed method. One completion of the segmentation task consisted of starting the ACM algorithm, interacting with the shape as the ACM iterates (possibly pausing the algorithm whenever necessary), and pushing the finish button of the space mouse when satisfied with the current ACM shape. Each participant performed the experiment in one session lasting about 40 minutes. The session was broken up into two force feedback conditions, that is, either force feedback was switched on or off. For each force feedback condition, the user conducted two segmentations on the two data sets earlier discussed, with conditions alternated between participants Before starting the experiment, each user was introduced to the system using a practice data set. After the experiment, the users were asked to fill in a questionnaire in order to provide us with subjective feedback.

6. Results and Discussion

6.1. Objective Results. The objective results gathered from the experiments were the time it took to perform the segmentations as well as the quality of the resulting segmentation. Time was measured from the moment the users pushed the start button until the moment they pushed finish button. It is important to mention that the users were allowed to first

navigate towards the ACM shape and inspect it for a while before starting the segmentation.

6.1.1. Segmentation Times. On average, participants took 206 seconds (3 minutes and 26 seconds) to complete one segmentation task. For the condition in which force feedback was present, the users performed the segmentations in 213 seconds and for the other condition participants were slightly faster, with 200 seconds. Although this difference can be considered negligible, a possible explanation for this is the fact that users could exert stronger forces when force feedback was switched off.

6.1.2. Segmentation Quality. To compare the quality between the traditional ACM and the version using haptics, we measured the overlap, $\mathcal{O} \in [0, 1]$, between the segmentations obtained with the two methods and a reference manual segmentation obtained with semiautomatic region growing. In terms of true and false positives and negatives, the overlap was quantified in 3 different ways:

Sensitivity (true positive rate): $\mathcal{O} = \text{TP}/(\text{TP} + \text{FN})$,

Dice similarity [22]: $\mathcal{O} = 2\text{TP}/(2\text{TP} + \text{FP} + \text{FN})$,

Jaccard similarity [23]: $\mathcal{O} = \text{TP}/(\text{TP} + \text{FP} + \text{FN})$.

Table 1 shows results obtained with the traditional ACM and the version using haptics, the latter further subdivided into force feedback switched on and off, as explained earlier. The results for the version using haptics are an average of the several runs of the application with the four participants of the experiments. In all cases, the traditional ACM performed slightly better, but the difference was in reality rather marginal. What is worth noticing is that, with the phantom stenosis, haptics with forced feedback switched on was noticeably better than its switched-off counterpart. We believe that this can be explained by the fact that the user can really feel the external force exerted by the edges of the oesophagus on the segmentation, therefore being able to counteract it with the aid of the haptics device.

6.2. Subjective Results. From the subjective feedback given by the participants, we were able to deduce several opinions which are interesting to take into account when creating segmentation systems like the one presented here.

We asked the participants to grade several statements using a Likert scale with 1 indicating strongly disagree and 5 strongly agree. An overview of the answers given about segmentation can be seen in Figure 7 and with regard to the
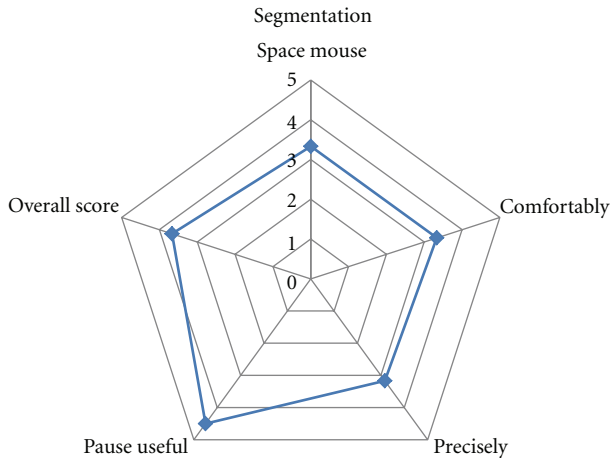
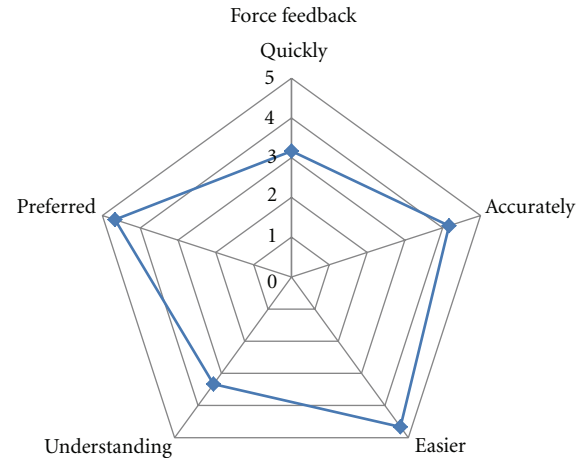Figure 7: Segmentation questionnaire scores from the subjective feedback.

Figure 8: Force feedback questionnaire scores from the subjective feedback.

force feedback in **Figure 8**. The values are the averages of the users' scores.

The scores for the segmentation indicate that the space mouse was perceived as an adequate device to perform navigation and system control with a slightly positive score ($M = 3.33$ SD $= 0.81$). Similarly, the mechanism was found to be comfortable and precise. A very interesting result is that the pause feature is found to be very useful ($M = 4.5$ SD $= 0.55$). Typically, users first let the ACM algorithm run for a while until they could see areas which were not deforming correctly. When those areas were found, they paused the simulation and edited the ACM shape until either satisfied with the end result or until they were reassured that the ACM algorithm could converge to the desired solution on its own. When asked to give the system an overall score, it was perceived as above average ($M = 3.67$ SD $= 0.52$).

The force feedback scores were in general very positive. A slightly worse score can be seen when the participants were asked if force feedback aids in performing the segmentation quickly. They were indeed indecisive about this (SD $= 1.33$ with $M = 3.17$). Some found it to be slower due to the fact that when force feedback was switched off they could push harder than their own comfort would allow when compared to when force feedback was switched on. Other users found force feedback reassuring and that it gave them a perceived increase in speed. Still, force feedback was overall found to be easier, preferred, and more accurate.

In postexperiment discussions, users reported it took some time to get acquainted with the system and the segmentation task. They also reported that without force feedback they felt less in control over the segmentation. One user even reported that it was impossible to carry out the segmentation without force feedback, since it was very hard to indicate if he or she was touching the shape or not. Although they also indicated that without force feedback it was possible to deform the shape more quickly, they preferred the higher sense of control offered when the force feedback was switched on rather than being faster when it was switched off.

Finally, the participants commented that the pause feature of the application's interface indeed gave them a mechanism to further control the segmentation process. It removed the need to compete with the deformations caused by the ACM algorithm, thus enabling easy manipulation of the shape. In the end, they found the system helpful in catching occasional mistakes made by the automatic part of the application.

6.3. Discussion. From the experimental results, we can deduce that our system shows potential to be used for segmentation of tracheal stenosis. It took less than 5 minutes to perform a segmentation in which the ACM was executed together with a user interacting with the ACM shape. The addition of force feedback was evaluated as beneficial to the quality of the interaction with the system. Furthermore, the ability to pause the ACM was found to be a very important feature which allows the user to take the time to apply his own changes to the current shape as well as to give full control over the final shape.

One of the observed downsides, however, is that it is very hard for a user to evaluate whether the current shape is good or bad. Sometimes it is very easy to diagnose wrong deformations, but with stenosis unexpected deformations are common. The visualisation used for the slices provides a good overview during interaction with the shape, but in order to diagnose the quality of the current shape the borders would have to be followed closely, which is not easy nor efficient. It is even potentially dangerous as accidental deformations can occur when unintentionally touching the shape borders while tracing them.

7. Conclusion

In this paper, we presented a 3D segmentation interface for tracheal stenosis from chest CT scans. We proposed a semiautomatic method that overcomes typical segmentation difficulties such as noise and the presence of neighbour organs in the image. Our method is based on Active Contour

Models augmented with 3D input from a haptic device. An evaluation was performed using two CT data sets and four users. The results indicated that the addition of haptics provided the users with valuable force feedback about the 3D surface and its deformation. Furthermore, the addition of a pause function during the interaction with the ACM proved to be an important part of the proposed method.

In future work, we would like to evaluate our system with more data sets and with users which are familiar with diagnosing tracheal stenosis, such that we could verify that our method is able to improve the diagnosing tasks.

References

[1] G. Qiu and P. C. Yuen, "Interactive imaging and vision—ideas, algorithms and applications," *Pattern Recognition*, vol. 43, no. 2, pp. 431–433, 2010.

[2] S. D. Olabarriaga and A. W. M. Smeulders, "Interaction in the segmentation of medical images: a survey," *Medical Image Analysis*, vol. 5, no. 2, pp. 127–142, 2001.

[3] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: active contour models," *International Journal of Computer Vision*, vol. 1, no. 4, pp. 321–331, 1988.

[4] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis: a survey," *Medical Image Analysis*, vol. 1, no. 2, pp. 91–108, 1996.

[5] N. Spittle and A. McCluskey, "Lesson of the week: tracheal stenosis after intubation," *British Medical Journal*, vol. 321, no. 7267, pp. 1000–1002, 2000.

[6] P. M. Boiselle, J. Catena, A. Ernst, and D. A. Lynch, "Tracheobronchial stenoses," in *CT of the Airways*, pp. 121–149, Humana Press-Springer, 2008.

[7] R. Pinho, K. G. Tournoy, and J. Sijbers, "Assessment and stenting of tracheal stenosis using deformable shape models," *Medical Image Analysis*, vol. 15, no. 2, pp. 250–266, 2011.

[8] J. M. Triglia, B. Nazarian, I. Sudre-Levillain, S. Marciano, G. Moulin, and A. Giovanni, "Virtual laryngotracheal endoscopy based on geometric surface modeling using spiral computed tomography data," *Annals of Otology, Rhinology and Laryngology*, vol. 111, no. 1, pp. 36–43, 2002.

[9] Y. Kang, K. Engelke, and W. A. Kalender, "Interactive 3D editing tools for image segmentation," *Medical Image Analysis*, vol. 8, no. 1, pp. 35–46, 2004.

[10] K. McGuinness and N. E. O'Connor, "A comparative evaluation of interactive segmentation algorithms," *Pattern Recognition*, vol. 43, no. 2, pp. 434–444, 2010.

[11] F. Heckel, O. Konrad, H. Karl Hahn, and H.-O. Peitgen, "Interactive 3D medical image segmentation with energy-minimizing implicit functions," *Computers and Graphics*, vol. 35, no. 2, pp. 275–287, 2011.

[12] A. Bornik, R. Beichel, and D. Schmalstieg, "Interactive editing of segmented volumetric datasets in a hybrid 2D/3D virtual environment," in *Proceedings of the 13th ACM Symposium Virtual Reality Software and Technology (VRST '06)*, pp. 197–206, November 2006.

[13] E. V. Zudilova-Seinstra, P. J. H. de Koning, A. Suinesiaputra et al., "Evaluation of 2D and 3D glove input applied to medical image analysis," *International Journal of Human Computer Studies*, vol. 68, no. 6, pp. 355–369, 2010.

[14] E. Vidholm, X. Tizon, I. Nyström, and E. Bengtsson, "Haptic guided seeding of mra images for semi-automatic segmentation," in *Proceedings of the 2nd IEEE International Symposium on Biomedical Imaging*, pp. 288–291, April 2004.

[15] E. Vidholm, S. Nilsson, and I. Nyström, "Fast and robust semi-automatic liver segmentation with haptic interaction," in *Proceedings of the MICCAI*, vol. 2, pp. 774–781, 2006.

[16] F. Malmberg, E. Vidholm, and I. Nyström, "A 3D live-wire segmentation method for volume images using haptic interaction," in *Discrete Geometry for Computer Imagery*, vol. 4245 of *Lecture Notes in Computer Science*, pp. 663–673, 2006.

[17] E. Vidholm, M. Golubovic, S. Nilsson, and I. Nyström, "Accurate and reproducible semi-automatic liver segmentation using haptic interaction," in *Medical Imaging: Visualization, Image-Guided Procedures, and Modeling*, vol. 6918 of *Proceedings of SPIE*, p. 69182Q, 2008.

[18] M. Harders and G. Székely, "Improving medical segmentation with haptic interaction," in *Proceedings of the Virtual Reality*, pp. 243–250, March 2002.

[19] G. C. Burdea, *Force and Touch Feedback for Virtual Reality*, Wiley Interscience, 1996.

[20] D. C. Ruspini, K. Kolarov, and O. Khatib, "The haptic display of complex graphical environments," in *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 345–352, 1997.

[21] A. Kulik, J. Hochstrate, A. Kunert, and B Froehlich., "The influence of input device characteristics on spatial perception in desktop-based 3D applications," in *Proceedings of the Symposium on 3D User Interfaces*, pp. 59–66, 2009.

[22] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945.

[23] P. N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*, Addison-Wesley, 2006.