A Knowledge-Based Approach to Automatic Detection of the Spinal Cord in CT Images

Neculai Archip*, Pierre-Jean Erard, Michael Egmont-Petersen, Jean-Marie Haefliger, and Jean-Francois Germond

***Abstract*—Accurate planning of radiation therapy entails the definition of treatment volumes and a clear delimitation of normal tissue of which unnecessary exposure should be prevented. The spinal cord is a radiosensitive organ, which should be precisely identified because an overexposure to radiation may lead to undesired complications for the patient such as neuronal disfunction or paralysis. In this paper, a knowledge-based approach to identifying the spinal cord in computed tomography images of the thorax is presented. The approach relies on a knowledge-base which consists of a so-called anatomical structures map (ASM) and a task-oriented architecture called the plan solver. The ASM contains a frame-like knowledge representation of the macro-anatomy in the human thorax. The plan solver is responsible for determining the position, orientation and size of the structures of interest to radiation therapy. The plan solver relies on a number of image processing operators. Some are so-called atomic (e.g., thresholding and snakes) whereas others are composite. The whole system has been implemented on a standard PC. Experiments performed on the image material from 23 patients show that the approach results in a reliable recognition of the spinal cord (92% accuracy) and the spinal canal (85% accuracy). The lamina is more problematic to locate correctly (accuracy 72%). The position of the outer thorax is always determined correctly.**

***Index Terms*—Image interpretation, knowledge representation, medical imaging, radiotherapy, spinal cord.**

I. Introduction

RADIOTHERAPY is an important ingredient in the often complex treatment procedures that are initiated in other to suppress different kinds of malignant tumors. The purpose of radiation therapy is to eradicate the tumor while minimizing the damage it causes to the surrounding healthy tissues. The spinal cord is an extremely radiosensitive, vital organ which should be spared as much as possible. A certain amount of exposure to radiation can induce a number of undesired neurological complications in the spinal cord (e.g., paralysis) [1]–[3]. The risk of such serious complications implies much smaller dosage tolerances for the spinal cord than for the tumor.

Successful radiotherapy relies on a precise planning and a thorough implementation of the radiation procedure. One of the problems of radiotherapy planning addressed here, is that it requires that the different tissues of interest, including the tumor and the surrounding (vital) organs, are located with a high accuracy. At present, radiation therapy is being planned by a radiologist and a radiotherapist in concert, based on a careful analysis of a computed tomography (CT) scan that covers the tumor and the surrounding tissues. The current planning procedure entails manual delineation of the spinal cord in each separate slice followed by an automatic reconstruction performed by the image analysis workstation that is connected with the CT-scanning device. Despite the existence of several semi-automatic approaches for planning of repetitive radiotherapy [4], [5], automatic detection of the spinal cord in CT images remains an unresolved problem. A factor that complicates the analysis further is the occasional presence of the spine around the spinal canal (Fig. 6).

Tumors are heterogeneous lesions, which exhibit growth patterns that are unique to each patient. As a consequence, the CT images cannot be acquired according to a standardized protocol but are subject to much inter-patient variation, e.g., compared with mammograms [6] or standard thorax radiographs [7] that are acquired in large numbers in Europe and North America. This rather high amount of variation in our image material impedes the application of a standard low-level image processing technique. For an image processing algorithm to be successful in our application, it should be *flexible* and also *transparent* to the radiologist and radiotherapist. A flexible approach can better cope with a high amount of inter-patient variation. Transparency of the image processing algorithms ensures that the experts can take over the image analysis, in case the automatic approach fails to give the desired result. To cope with the requirements of flexibility and transparency, we present a knowledge-based approach to automatic image analysis. The basic components of our system are the so-called *anatomical structures map* (ASM) [8] and the *plan solver*, a task-oriented module that controls the sequence in which the subtasks are performed. The ASM and the plan solver are designed such that they capture parts of the anatomical and procedural knowledge that is currently being used for manual image interpretation.

This paper is structured as follows. First, existing approaches to knowledge-based image interpretation are discussed. Then, we consider different archetypes of knowledge that are presently used to solve the spinal cord detection problem. Subsequently, we give a detailed description of the knowledge-based archi-

Manuscript received May 30, 2002; revised September 12, 2002. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was M. W. Vannier. *Asterisk indicates corresponding author*.

*N. Archip was with the Computer Science Department, University of Neuchâtel, Emile-Argand 11, CH 2007, Neuchâtel, Switzerland. He is now with the Department of Electrical and Computer Engineering, The University of British Columbia, 2356 Main Mall, Vancouver, BC V6T 1Z4, Canada (e-mail: neculai.archip@ece.ubc.ca).

P.-J. Erard is with the Computer Science Department, University of Neuchâtel, Emile-Argand 11, CH 2007 Neuchâtel, Switzerland.

M. Egmont-Petersen is with Institute of Information and Computing Sciences, Utrecht University, Centrumgebouw Noord, De Uithof 3584CH Utrecht, The Netherlands.

J.-M. Haefliger and J.-F. Germond are with La Chaux-de-Fonds Hospital, CH 2300 La Chaux-de-Fonds, Switzerland.

Digital Object Identifier 10.1109/TMI.2002.806578

(a)

(b)

Fig. 1. The two types of slices. (a) Spinal canal completely surrounded by bone. (b) Spinal canal partially surrounded by bone.

tecture consisting of the ASM and the plan solver. Following this description, the low-level (atomic) image operators are described in detail. In the experimental section, we report the results obtained by applying our approach to the CT images obtained from 23 patients before they underwent radiation therapy. This paper ends with a discussion of the results and issues for future research.

II. Knowledge Guided Image Processing

The anatomical information present in our image material is very complex and hard to formalize in way that makes computer-based image interpretation feasible. As argued in the *Introduction* section, our image material is characterized by a large amount of inter-patient variation. This variation makes it difficult to develop standardized low-level image processing algorithms that make feasible an automatic detection of the spinal cord in CT images. Instead, we present a novel, knowledge-based top-down approach to image interpretation. Our approach was originally inspired by the manner in which a radiologist and radiotherapist interpret the CT images before the actual radiotherapy is planned.

A. Image Acquisition and Interpretation

In clinical routine, the radiotherapist performs a request containing the questions which should be resolved by the radiological examination. Examples of questions are: where is the tumor located? how far is it from the spine? are there other healthy tissues that will be exposed to radiation? etc.

The image acquisition is performed according to a standard protocol, which contains general guidelines for how CT images should be obtained for planning of radiotherapy. The details of an acquisition are chosen such that the tumor of the particular patient is visualized in the best possible way. In general, a number of aspects should be taken into account in order to acquire CT images in such a way that the relevant findings can be established. Wegener [9] points out that there is a strong relationship between what region, organ or lesion is examined and how the image should be acquired, including imaging parameters (slice thickness, slice interval, scanning time), and contrast administration (presence/type of contrast agent, injection speed, concentration).

After the CT images have been acquired, the interpretation is performed by a radiologist and a radiotherapist in concert. The image assessment relies on both morphological and densitometric findings. Grimnes mentions a number of general aspects that influence the interpretation of CT images [10]:

- — the typical size and shape of the objects (bones, organs, and other tissues);
- — the variation in size and shape of the objects;
- — the expected Hounsfield unit (HU) value range associated with each tissue;
- — the variation in the HU value range associated with each tissue;
- — typical response of an organ to the contrast tracer that is used;
- — organs and blood may change their expected HU range in light of disease;
- — biological variation;
- — and social context.

The radiological analysis results in a *synthesis* of the clinically relevant findings present the CT images, while taking the abovementioned aspects into account. The ultimate goal of any computer system for image interpretation should be to produce such an image synthesis, either automatically or in an interactive manner, e.g., through a dialogue with the radiologist.

B. Existing Approaches to Knowledge-Based Image Interpretation

The literature on computer-based image interpretation describes a large number of architectures, systems and approaches. Among the conventional approaches for image interpretation, some focus on architectural aspects of the scene (the spatial configuration composed by the objects that are present); in other approaches an extensive knowledge base and an advanced reasoning strategy form the major components [11]–[14]. Also

TABLE I
RELATIONSHIPS BETWEEN STRUCTURES

Relation	Description
isInside	Object1 *isInside* Object2 ⇔ allmost all the voxels of Object1 are included in Object2
isSorounded	Object1 *isSurrounded* Object2 ⇔ allmost all the voxels of Object1 are included in Object2 and Object1 is the only object which respect Object1 *isInside* Object2
isAtMedialAxis	Object1 *isAtMedialAxis* Object2 ⇔ the center of Object1 is approximately the same as the center of Object2
isNeighbor	Object1 *isNeighbor* Object2 ⇔ the Object1 has a common border with Object2
isAtLeft	Object1 *isAtLeft* Object2 ⇔ allmost all the voxels of Object1 are at the left border of Object2
isAtRight	Object1 *isAtRight* Object2 ⇔ allmost all the voxels of Object1 are at the right border of Object2

probabilistic systems were developed for knowledge-guided image interpretation [15]–[17]. Several blackboard and other knowledge-based systems were developed specifically for interpretation of medical images: The ERNEST system has been developed for interpretation of scintigraphic images and magnetic resonance (MR) images [18]. The system VIA-RAD [19] applies four diagnostic strategies, obtained from the radiological domain, to perform image interpretation. Brown *et al.* [20] present a knowledge-based system for lung detection in CT images. A system for object recognition in CT images of the brain is presented in [21]. An architecture has been developed for interpretation of abdominal CT images [22]. A task-based architecture to interpretation of MR images of the brain is introduced by Gong *et al.* [23].

In computer-based systems for interpretation of medical images, one or more of the following archetypes of knowledge may be modeled [24]:

—*structural knowledge*, which can contain information about the physical world (human anatomy, e.g., normal structures such as lungs, spinal canal, lamina, spinal cord, thorax, etc.);

—*dynamic knowledge*, which can contain information about possible normal and abnormal processes (human physiology and pathology);

—*procedural knowledge*, which divides a request (e.g., image synthesis) into a sequence of subtasks that can be performed by specific image processing algorithms.

In some applications, a satisfactory image synthesis can be obtained from solely one type of knowledge. For example, in perfusion analysis of bone tumors dynamic knowledge is sufficient for making a distinction between viable tumor and necrosis [25]. In other applications, all three types of knowledge may be a prerequisite for a successful image synthesis. Spinal cord detection and subsequent planning of radiotherapy rely primarily on structural knowledge components: where is the tumor located, the spine, etc., and on the procedural knowledge that is needed to describe how the CT images should be analyzed [26], [27].

C. Knowledge Representation in Medical Image Analysis

We will present an approach for semi-automatic image interpretation that uses a knowledge base to link different low-level image processing algorithms to a particular request. For a solution of the problem addressed—spinal cord detection—a combination of *structural* and *procedural* knowledge suffice, because the pathologic growth process of the tumor does not have to be taken into account. This demarcation implies that our knowledge base should contain medical knowledge about organs and possible pathologic structures, i.e., components of the tumor. The knowledge-base is used to guide the image interpretation but also to specify the parameters of the basic algorithms. The architecture presented here is inspired by *frame* systems [28]. Each anatomical structure is represented as a prototype, and its properties as *slots*, which may obtain values by attached *demons*.

Structural Knowledge: The core of our system is the so-called ASM, which was presented earlier in [8]. A set of properties (related to shape, position, densitometric ranges) is used to characterize each of the normal structures, the organs, bones and the vascular system, that are represented in the ASM. The spatial arrangement of these objects is represented as a semantic network. A very simple grammar was also introduced that makes it feasible to express the semantic relations that pertain to our application, see Table I.

Procedural Knowledge: The structural knowledge base is merged with a task oriented architecture, the *plan solver*, which contains the procedural knowledge that is needed to perform the image interpretation. The involved clinicians make use of so-called reference objects (e.g., body or lamina) to direct their

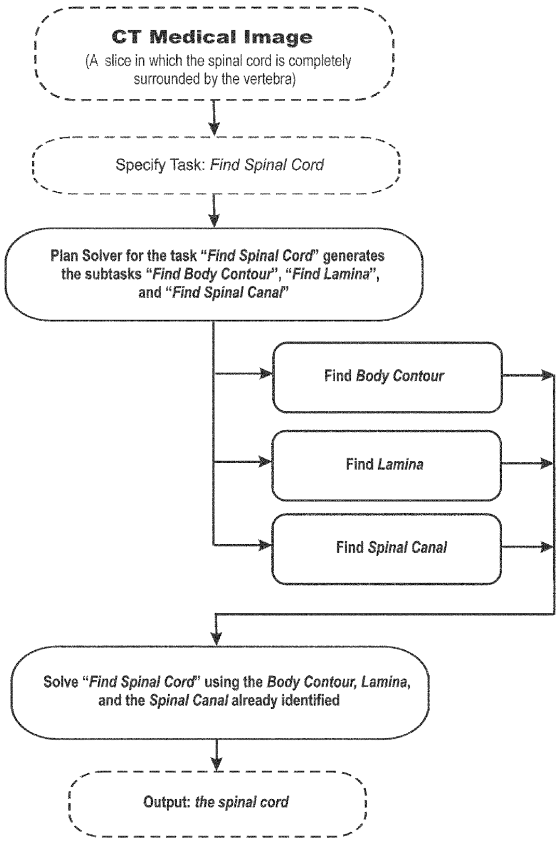

Fig. 2. Plan for a task: general architecture in the case of spinal cord identification.

focus of attention. Although the architecture of the plan solver was originally inspired by the approach followed by the involved clinicians, the task-based structure also makes it possible to recognize and locate complex objects while benefiting from more simple (basic) object detections. Algorithms developed for the recognition of complex objects use so-called *reference objects* to set their initial configuration or constrain the final solution.

The task oriented architecture is responsible for running the plan solver, which dispatches a task, e.g., detect spinal canal, into subtasks [23]. Which subtask should be dispatched, depends *the reference objects*. $Object_x$ is reference object for $Object_y$ if:

- there is a direct, spatial relation between $Object_x$ and $Object_y$ (e.g., *isNeighbor*, *isInside*), and
- $Object_x$ has an segmentation algorithm that does not depend on $Object_y$. Hence, $Object_x$ can be detected without any knowledge of $Object_y$.

When the plan-solver is called with the request *Find* $Object_x$, it identifies the subtasks that should be performed in order to fulfill the request, i.e., which objects are reference objects to $Object_x$. The list with reference objects found is the list with the subtasks to be performed. The plan solver module relies on a global positioning system (along the axes x, y, z), (Fig. 5) which maps each of the detected organs to world coordinates.

III. Anatomical Structures Map

The ASM establishes a number of spatial relations between the objects that are typically distinguished in the CT images used for planning of radiotherapy in our clinic. The architecture of the ASM lends its inspiration from frame systems, a well-known concept in the artificial intelligence literature. We chose to represent the anatomical information in two–dimensional (2-D) slices. More specifically, the ASM represents spatial relations between the objects (e.g., spine, lamina, and tumor) as well as the general category of each object: bone, air and tissues (see Fig. 3). We discern these particular categories of tissues for the following reasons. Objects belonging to the first two categories have either a very high or a very low HU level (bones versus, e.g., the air compartment of a lung). For these two types of objects, a threshold-based technique is in most cases sufficient for a reliable segmentation result. Tissues (e.g., organs), on the other hand, cannot be identified by thresholding within a specific HU range. For objects belonging to this third category, a reliable segmentation needs to be based on two kinds of information: the locations of the already detected reference objects and the results of texture segmentation.

The main object represented is the *body contour*, which comprises other organs. It has a so-called *independent* segmentation scheme as it is possible to detect the body by a basic image processing algorithm, in this case by thresholding (see Section V-A-1).

The structures that are more difficult to segment include the spine, the lamina and the spinal canal. The spine contains mainly bone so it has a very high HU range, and thresholding is used to detect it. All the subparts of the spine consist of mainly bone cortex so a threshold method is used to detect these objects. The spinal canal consists mainly of tissue but is completely surrounded by the spine, i.e., *Spinal Canal* IsInside *Spine*. We use a region growing scheme to segment it, chiefly because the border of the spinal canal has a high contrast compared with the surrounding bone (difference HU bone-tissue), see Section V-A-3.

Finally, we represent the lung information, the ribs and any lung tumors. Our approach to lung tumor detection is described in [29].

IV. The Plan Solver

The ASM is used as aid when partitioning a request (e.g., locate spine) into subtasks and further into atomic image processing tasks that are performed by dedicated routines. This hierarchical partitioning takes place in the *plan solver* module, which links the spatial relations in the ASM with atomic image processing algorithms. The plan solver uses an *inheritance scheme* to determine the appropriate segmentation approach for a particular object or tissue (see Fig. 4). An object connected with another (reference) object by an **isA** relationship inherits the segmentation method of that object.

We make a distinction between different types of atomic segmentation methods that are used for object recognition in our application: the threshold-based methods (for the bones in this case, but also for the lungs in relation to lung tumors) and region based methods (for the spinal canal in this case).

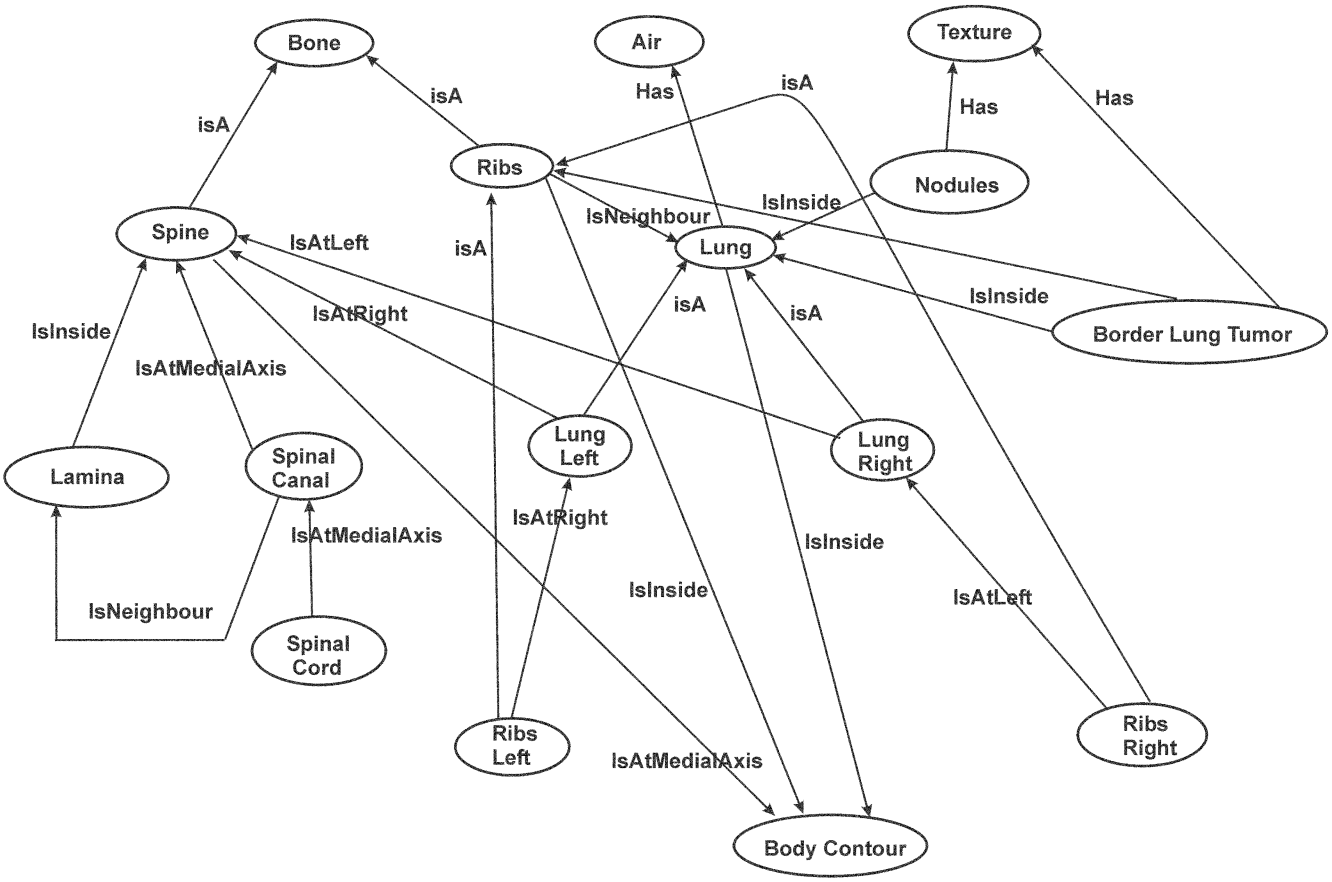

Fig. 3. Relationships between structures, the *ASM*.

For the threshold-based methods, it is important to restrict the area to which they are applied. This is accomplished by using so-called reference objects. Reference objects are specified by the following relations in the ASM: **isAtLeft**, **isAtRight**, **isInside** and **isSurrounded**, which are applied in a recursive top-down detection procedure.

For region based approaches, the reference objects are found between the objects with the relationship **isNeighbor** or **isVerticalAxis**. When a certain (sub)request *Find Object*$_x$ is dispatched, the plan solver tries to fulfill the request by choosing the appropriate segmentation methods. These are either specified directly (for certain organs like the spinal canal, which is detected by region growing), or indirectly by inheritance from the reference objects by the relationship **isA**. Depending on the applicability of the chosen segmentation method on the particular image slice, the reference objects are located successfully.

We illustrate the functionality of the plan solver by two example requests: *Find Lamina* and *Find Spinal Canal*. The first object, *Lamina*, does not have its own dedicated segmentation methods (no demon present) so Lamina is found by the inheritance structure based on the link **isA**. *Lamina* **isA** *Spine*, which also does not have its own dedicated segmentation method. Finally, *Lamina* **isA** *Bone* which has a thresholding segmentation method attached. As *Lamina* is connected by the link **isA** to *Bone* via *Spine*, *Lamina* is segmented by thresholding. The inference mechanism proceeds by looking for the objects linked to *Lamina* by the relation **isInside**. The only object where *Lamina* is inside, is the body (body contour), which this way becomes a reference object for *Lamina*. So the task *Find Lamina* has as subtasks: *Find Body Contour* and *thresholding*, the latter takes place only inside the *Spine*.

In the second example, *Find Spinal Canal*, a dedicated segmentation method is specified: region based segmentation. So we are looking for the objects which could give us a starting point for the region growing algorithm. Thus, we are looking for the objects connected with relationships **isNeigbor**, **isAtVerticalAxis**, which are *Body Contour* and *Lamina*. *Body contour* has its own segmentation scheme, which is why it is a reference object for the *Spinal Canal*. The *Lamina*, as it is presented earlier, has as reference object *Body Contour*, which does not involve the *Spinal Canal*. So the *Lamina* is the second subtask for the task *Find Spinal Canal*.

V. Spinal Cord Three-Dimensional (3–D) Detection

The 3-D image interpretation method presented in this paper was developed in an attempt to model certain aspects of the knowledge that is used for human interpretation of CT images of the thorax. An organ of interest is segmented by identifying in each slice the contours of interest, and by using information obtained from adjacent slices to improve the result further. For the spinal cord, the occasional presence of spine around the spinal

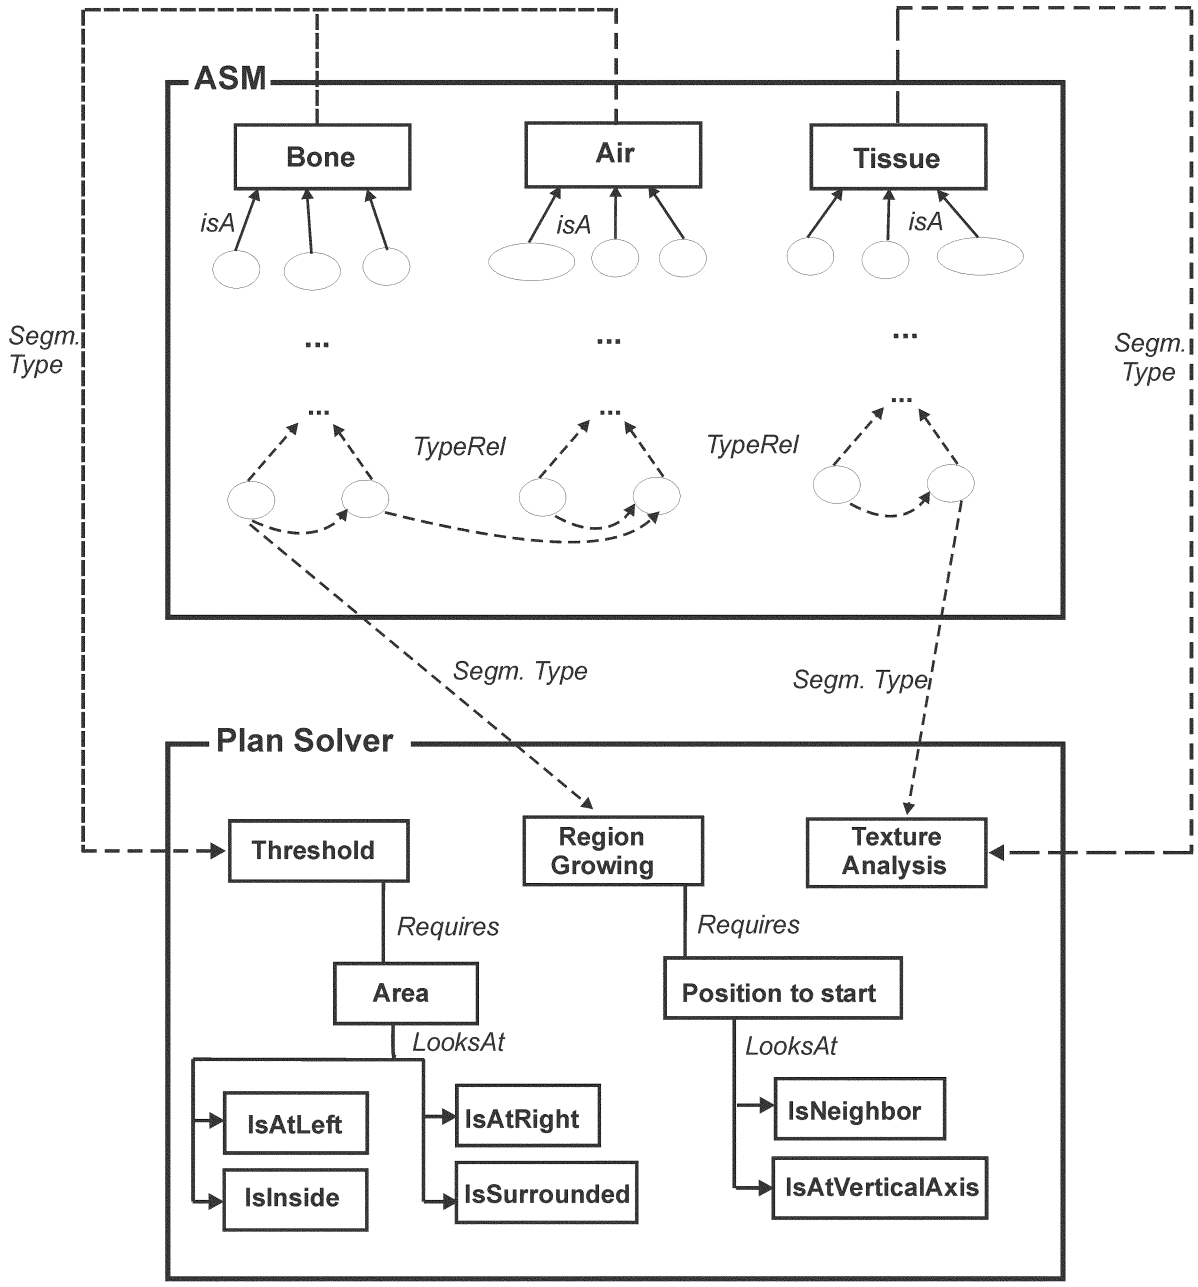

Fig. 4. Relationships between the ASM and the *plan solver*. All relationships between structures are presented, and on the *plan solver* side segmentation methods as well as parameters for each of it. Also the connections between structures and corresponding segmentation methods are presented.

canal complicates the delineation of its contour. Moreover, the same segmentation scheme cannot be used in all the slices. In this section, we first present the 2-D segmentation of spinal cord, which is based on the ASM and the plan solver that are applied to the slices in which the spinal canal is completely surrounded by spine. Subsequently, the procedure responsible for detection of the 3-D spinal canal contour is described. Finally, the methods used in the case of failure of the standard procedures (in the slices where the spinal canal is not surrounded by spine) are presented.

A. Two-Dimensional (2-D) Spinal Cord Detection Based on the ASM

For the task of identifying the spinal cord contour in a slice, the plan solver is dispatched. Its subtasks rely on information from the ASM. Fig. 2 illustrates how the spinal cord is being detected by our knowledge-based approach. The structures that aid the detection of the spinal cord are body contour, a region of the spine (called lamina), and the spinal canal (see also Fig. 5).

1) Body Contour Identification: The transition between the body (contour 1 in Fig. 5) and the outside air is very strong,

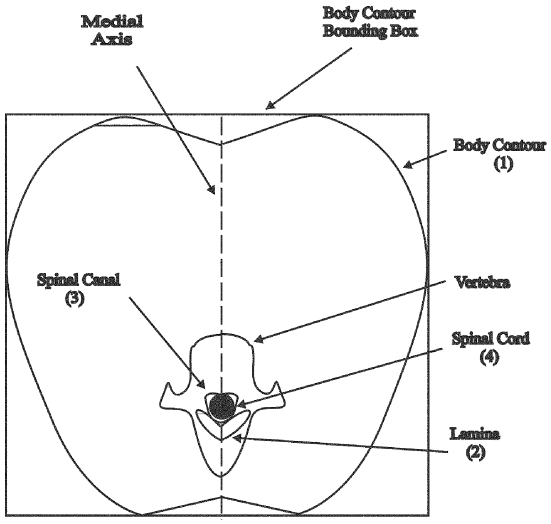

Fig. 5. Contours detected for spinal cord detection. First *Body contour* is detected, then *Lamina*, followed by *Spinal Canal* and finally *Spinal Cord*.

which makes it rather straight-forward to find the contour around the thorax. Moreover, the body is generally the only object in the image. The voxels with a gradient exceeding a threshold value ε are likely to form part of the border between body and air. Based on correlational analysis of the HU histograms of the body and air in a pilot study, the value of ε was found. The algorithm 1 is used to delineate the contour around the body.

Algorithm 1 Body Contour Identification

Require: Image I
Ensure: the abdomen contour
1: compute the gradient of the image using a Sobel-like operator;
2: in the middle column of the image, search the first pixel which has the gradient higher than a threshold ε;
3: this is the first point on the body contour;
4: starting from this point, follow in the clock-wise direction the high gradient, until it reaches the first point of the contour.

Because of its importance (all the other structures are *Inside* the body contour), the body contour identification is a subtask which is performed to accomplish each other request.

2) Lamina Identification: The lamina contour (contour 2 in Fig. 5) uses body contour as a reference object. In the ASM, the segmentation scheme associated with the lamina is a threshold operation. The threshold operator is applied to the voxels that occur inside the body contour. The lamina has a very high HU range (650–1200 HU). This is not the only structure with such a high intensity range. Other structures like the *sternum* and *scapula* might also be detected by application of a threshold operator. By restricting the threshold operator to a smaller region

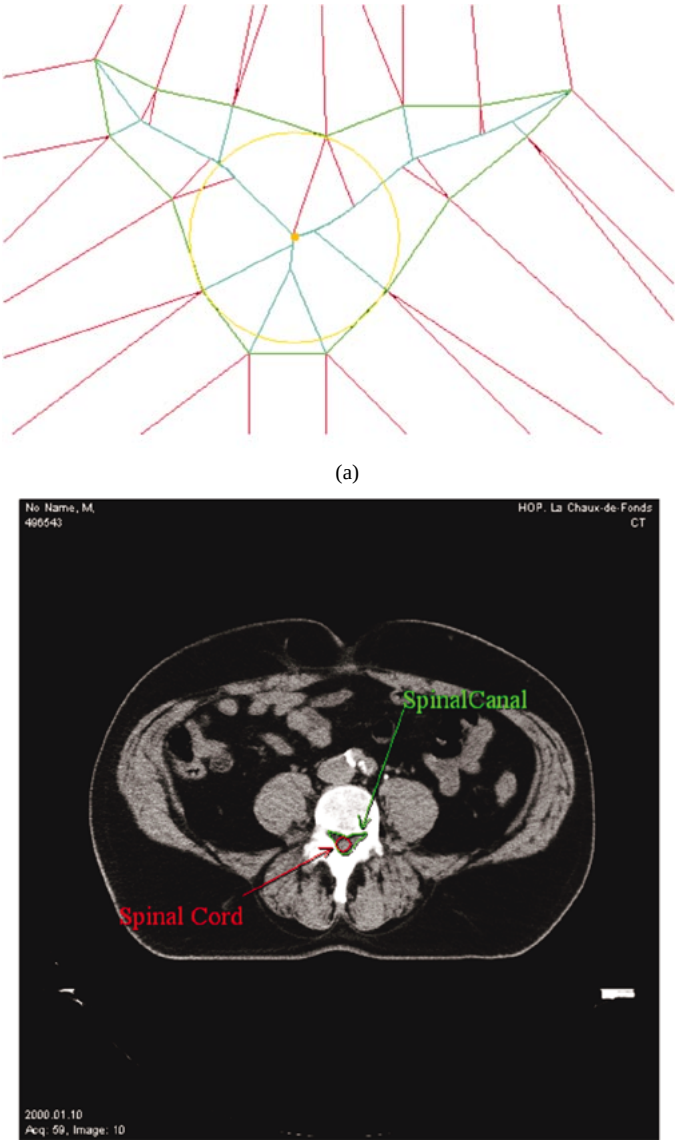

(a)

(b)

Fig. 6. MIC in the spinal canal polygon. (a) The Voronoi Diagrams used for detection of the medial axis, which gives the center of MIC. (b) This solution applied for spinal cord detection.

of the abdomen (centered at the medial axis), thresholding in most cases finds the lamina accurately.

3) Spinal Canal Detection in Two Dimensions: The contours of the lamina and thorax are used to detect the spinal canal (contour 3 in Fig. 5). There is a strong transition from lamina to the spine (large HU difference bone-tissue). Consequently, a region growing algorithm is used [30]. Two problems are related to the region growing algorithm. First, the homogeneity of the pixel intensities in the region may not be guaranteed. To cope with this problem, a histogram based method [31] is combined with the *a priori* knowledge about the typical HU range of the spinal canal. A pilot experiment has been performed to find the optimal range of HU values.

A second problem is how to set the seed point—the starting point of the region growing algorithm—automatically. This is accomplished by using the relative locations between body contour and lamina in relation to the spinal canal. More specifically,

the spinal canal and the spine have the same medial axis (represented by the relationship **isAtMedialAxis**). So, by detection of *lamina* (which is **isInside** the spine), the position of the seed point for the spinal canal is obtained (being on the medial axis and higher than the top limit of *lamina*).

4) Spinal Cord Detection in two dimensions: The problem of spinal cord detection (contour 4 in Fig. 5) reduces to finding the maximal inscribed circle in the polygon that represents the spinal canal (see also Fig. 6). The problem is solved by computing the skeleton of the polygon using an efficient algorithm [complexity $O(n \log n)$], which was presented in [32] and [33].

B. Spinal Canal Detection in Three Dimensions

The problem of 3-D spinal canal detection is based directly on the procedure for spinal canal detection in two dimensions presented in Section V-A-3. However, this scheme cannot be applied successfully to all slices because the spinal canal is not always surrounded by the spine. Instead, the algorithm for 3-D spinal canal detection first identifies the spinal canal each slice using the algorithm 2. The first step is to apply the 2-D algorithm presented in the previous section to identify the spinal canal in the first slice. It uses no information about whether the spinal canal is surrounded completely by bone. A procedure verifies (line 4) whether the spinal canal was identified correctly. This procedure uses information about the position, the intensity and the area of the region detected by the 2-D algorithm. If the algorithm failed to identify the spinal canal correctly in the first slice, the same 2-D algorithm is applied to the next slices (lines 4–8), until it succeeds in finding the contour of the spinal canal in as many slices as possible.

Once a contour around the spinal canal has been found, the algorithm uses it as a reference in the neighboring slices in two ways: First, it is used to verify the candidate contour for spinal canal in the adjacent slices (lines 14, 28; assuming a small difference between the contours of the spinal canal in two consecutive slices). The second way is to use a spinal canal contour as information to guide the segmentation scheme in the adjacent slice (in case the 2-D algorithm fails to identify the spinal canal correctly—lines 13, 27). These two applications of the spinal canal contour already identified are presented in the next section. In the kth slice, the contour of the spinal canal is detected (line 10). The 3-D algorithm proceeds from the $k-1 \rightarrow 1$ and $k+1 \rightarrow nrTotalSlices$ slices, applying the 2-D detection algorithm. In case of failure, it chooses one of the alternative methods presented in the next section. The evolution of the algorithm in two consecutive slices is illustrated in Fig. 7.

Algorithm 2 Spinal Canal 3-D Detection

```
Require: medical exam 3-D, nrTotalSlices
Ensure: the list with spinal canal contour identified in all the slices of the exam, ListSCContours
1: SetActiveSlice(1)
2: SCC ← DetectSpinalCanalContourUsingASM()
3: k ← 1
4: while NOT VerifyCandidateSpinalCanal(SCC) do
5:   k ← k + 1
6:   SetActiveSlice(k)
7:   SCC ← DetectSpinalCanalContourUsingASM()
8: end while
9: /*in the slice k, spinal canal is identified*/
10: ListSCContours(k) ← SCC
11: for j = k − 1 to 1 do
12:  SetActiveSlice(j)
13:  SCCNew ← DetectSpinalCanalContour(SCC)
14:  if NOT VerifyCandidateSpinalCanal(SCC, SCCNew) then
15:   /*spinal canal not correctly identified so use snakes*/
16   SCCNew ← ModifyUsingSnake(SCC)
17:  end if
18:  SCC ← SCCNew
19:  ListSCContours(j) ← SCC
20: end for
21: SCC ← ListSCContours(k)
22: for i = k + 1 to nrTotalSlices do
23:  SetActiveSlice (i)
24:  SCCNew ← DetectSpinalCanalContourUsingASM()
25:  if NOT VerifyCandidateSpinalCanal(SCCNew, SCC) then
26:   /*spinal canal not correctly identified with ASM*/
27:   SCCNew ← DetectSpinalCanalContour(SCC)
28:   if NOT VerifyCandidateSpinalCanal(SCC, SCCNew) then
29:    /*spinal canal not correctly identified so use again snakes*/
30:    SCCNew ← ModifyUsingSnake(SCC)
31:   end if
32:  end if
33:  SCC ← SCCNew
34:  ListSCContours(i) ← SCC
35: end for
```

Two procedures are used to check the results of the spinal cord detection algorithms. The first one (line 4), *VerifyCandidateSpinalCanal(SCC)*, uses specific *a priori* knowledge about the spinal canal region: *position*, *area*, *intensity*, and *shape*. If the properties of the candidate comply with the predefined parameters of our model, the region is recognized as *spinal canal*, otherwise it is rejected.

The second procedure (line 14) *VerifyCandidateSpinalCanal*(SCC, SCCNew) uses a contour obtained in an adjacent slice, against which it verifies the properties of the new contour detected in the current slice. If the differences between the two

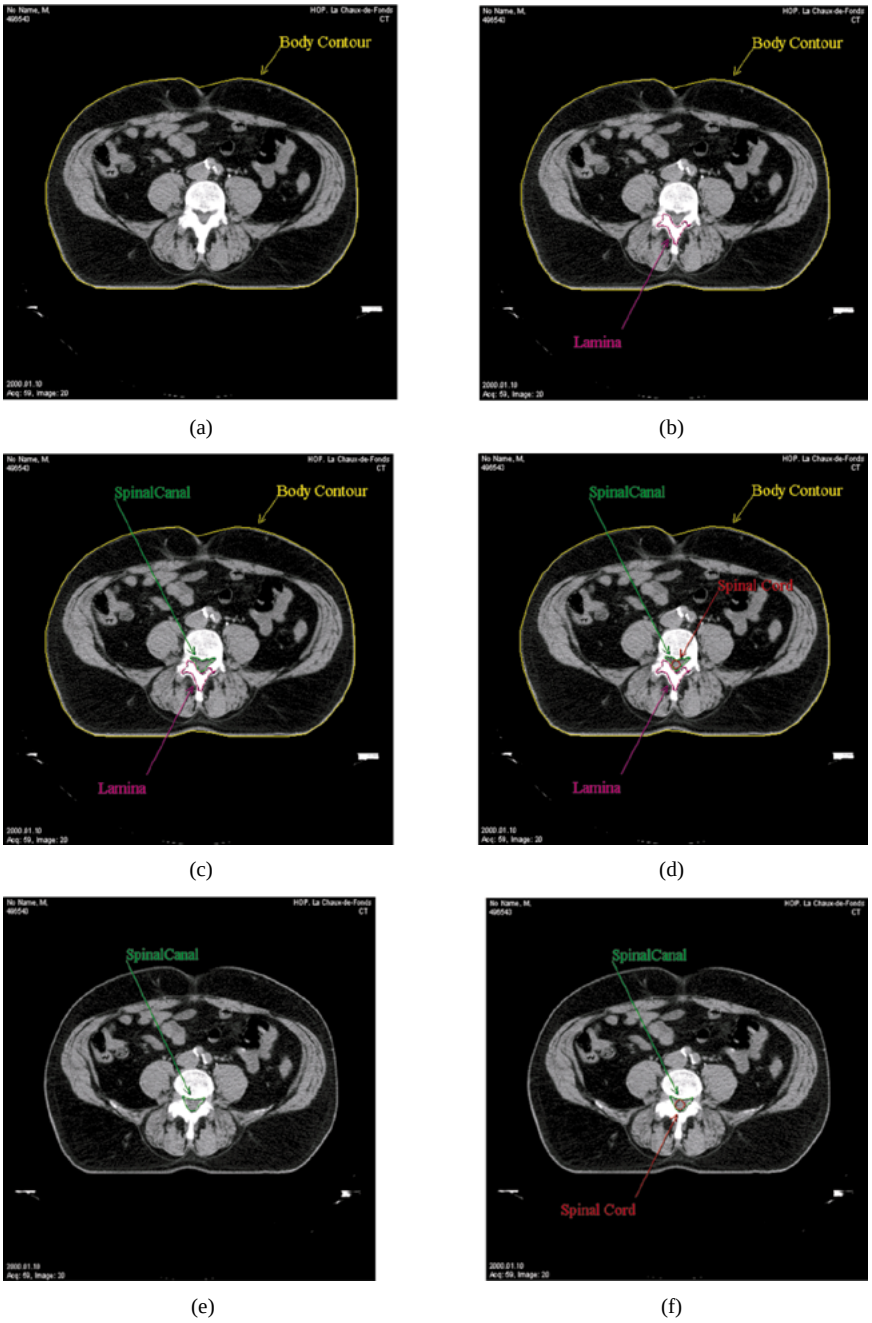

Fig. 7. Progress of our algorithm for spinal cord identification. (a) In the first slice identify the body contour. (b) Find a spine part. (c) Using the bone part identified, find the seed for the Region Growing, which identifies the spinal canal. (d) Apply MIC algorithm to find the spinal cord. (e) Propagate the spinal canal contour in the next slice and improve it using snakes. (f) Apply again MIC to find the spinal cord.

contours with respect to: *position*, *area*, *intensity*, and *shape*, are smaller than a set of prespecified ranges, the new candidate contour is recognized as being the *spinal canal*.

C. When Segmentation of the Spinal Canal Fails

In case the general scheme based on the default algorithms applied by the plan solver fails to detect an appropriate contour around the spinal canal, the system backtracks and uses either: a region based method or snakes. Both use the (already approved) contour around the spinal canal detected in an adjacent slice for initialization.

1) Finding the Spinal Canal by Region Growing: Occasionally, the general scheme for detection of the spinal canal fails because it cannot identify the lamina region, even when the spinal canal is completely surrounded by bone cortex. In these cases, a region based segmentation technique works well and is applied to the slice (lines 13, 27). The problem is to find the seed point for the region growing process. We use as seed point the center of gravity of the spinal canal region identified in an adjacent slice, thereby assuming spatial continuity of the spinal canal. To compute the center of gravity of a region given by a function f, statistical moments are used

$$m_{pq} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x^p y^q f(x,y)\,dx\,dy \quad \text{with } p,\ q = 0,\ 1,\ 2,\ \ldots \tag{1}$$

(in the continuous case) and

$$m_{pq} = \sum_i \sum_j i^p j^q f(i,j) \qquad \text{with } p,\, q = 0,\, 1,\, 2,\, \ldots$$

for the discrete case. Thus, the center of gravity is defined as

$$\overline{x} = \frac{m_{10}}{m_{00}} \quad \text{and} \quad \overline{y} = \frac{m_{01}}{m_{00}}.$$

2) The Use of Snakes: In case the general scheme fails to detect the spinal canal, e.g., because the spinal canal is not completely surrounded by the spine, the system uses a snake based method [34].

Given the spline $v(s) = (x(s),\, y(s))$, the energy function of the snake is defined as

$$E_{\text{total}} = \int_0^1 E_{\text{int}}(v(s)) + E_{\text{image}}(v(s)) + E_{\text{con}}(v(s))\, ds. \quad (2)$$

E_{int} represents the internal energy of the spline, composed of a first-order term controlled by $\alpha(s)$, which makes the snake act like a membrane, and the second-order term controlled by $\beta(s)$, making the snake to act like a thin plate

$$E_{\text{int}} = (\alpha(s)|v_s(s)|^2 + \beta(s)|v_{ss}(s)|^2)/2 \quad (3)$$

E_{image} is given by

$$E_{\text{image}} = -|\nabla I(x,\, y)|^2 \quad (4)$$

so that the snake is attracted by the contours with large gradients. Finally E_{con} incorporates the external forces specified by the user. The problem of initialization of the snake is solved using the result obtained in an adjacent slice. We chose a greedy strategy as in [35] to search for the best snake contour.

Combining these energy equations results in the snake algorithm 3. The gradient is hereby approximated by a Sobel operator [31].

Algorithm 3 Snake General

```
Require: Image I, snakes parameters α, γ,
  window size, contour v, nrPoints
Ensure: the contour v modified
1: G ← ComputeGradient(I)
2: ∀i = 1, nrPoints β(i) ← 1
3: finish ← false
4: while NOT finish do
5:  for i = 1 to nrPoints do
6:   ModifyPoint(i, newX, newY)
7:   v(i) ← (newX, newY)
8:   Evaluateβ()
9:  end for
10:  if ConditionsFinishOk() then
11:   finish ← true
12:  end if
13: end while
14: return v
```

VI. Experiments

We evaluated our knowledge-based approach by applying the system consisting of the ASM and the plan solver on 3-D CT datasets of 23 real patients, all scanned at the *La Chaux de Fonds* hospital in Switzerland. All 23 patients had a tumor present in the thorax. After the CT examination, each patient underwent radiotherapy in the hospital. Our population consisted of 9 male and 14 female patients. Their age varied from 37 to 79 years with a mean of 58 years and a median of 59 years. The number of CT slices per exam varied from 9 to 97 with a mean of 45 slices and a median of 38 slices. The images were obtained from a CT scanner from Picker and were acquired with a slice thickness of 3 mm and an inter-slice distance of 3 mm.

We evaluated our approach according to two criteria: *accuracy* and *computational cost*. The accuracy is defined as the relative number of acceptable contours of a particular type that can be detected in an exam. The computational cost is the total execution time (in seconds) required to find all contours of a particular type in an exam (one contour per slice). We distinguished four types of contour: Spinal cord, Spinal canal, Lamina, and (outer) Thorax. Whereas the computational time is straight-forward to compute, medical expertise is needed to assess the contours that were found by our system. A radiologist skilled in radiotherapy planning was asked to accept or reject each contour in each slice among all 23 patients. In our case, evaluation was performed using a visual inspection of the contours projected on the CT image slices. The radiologist decides for each of the contours obtained with our system whether it is located precisely or not. The maximal deviation around the spinal cord contour (measured as the distance perpendicular to the contour) accepted by the radiologists is 1 mm.

A. Accuracy

In Table II, results of the experiments on the real clinical data are shown. The accuracy is computed as the number of slices in the exam in which the particular contour was located correctly. In Exam 1, for example, 91.8% of the contours were located around the spinal cord with a sufficient precision. The average accuracy of the spinal cord contours among all patients is 91.7%, the average accuracy per slice lies within the range 80% to 100%. The spinal canal is more difficult to detect. The average detection accuracy among all patients is 85.3%, the average accuracy per slice lies within the range 60% to 100%. The lamina is the most difficult structure to detect for our approach. The average accuracy among the 23 patients is 72.1%, the range is 33%–100%. Finally, the thorax is located correctly in all slices among all 23 patients. The body is easy to detect because of the sharp transition from the surrounding air to human tissue.

Among the four structures, the accuracy of each contour is rather correlated between the spinal canal and spinal cord, 0.594, the correlation coefficients between the other types of contours are all below 0.15.

Due to the parameters used for the snakes algorithm which segments the spinal canal, the contours detected contain just false positive regions. Then, by the use of MIC algorithm (the maximum circle approximation for the spinal cord), occasionally the spinal cord contours contain false positive as well.

TABLE II
RESULTS FROM THE EVALUATION OF OUR SYSTEM USING THE CT IMAGES OF 23 PATIENTS

Medical	No. of	Age	Sex	Sp. Cord		Sp. Canal		Lamina		Thorax	
Exam ID	Slices	(years)		*Succ.*	*Texec*	*Succ.*	*Texec*	*Succ.*	*Texec*	*Succ.*	*Texec*
Exam 1	37	68	F	91.9%	65s	91.89%	62s	48.64%	42s	100%	34s
Exam 2	87	60	M	94.26%	153s	93.10%	148s	98.85%	92s	100%	73s
Exam 3	27	58	F	88.89%	40s	85.18%	39s	74.07%	25s	100%	18s
Exam 4	70	62	M	91.43%	108s	90%	105s	62.85%	68s	100%	52s
Exam 5	87	59	F	93.11%	177s	89.65%	173s	82.75%	112s	100%	90s
Exam 6	15	61	M	86.7%	31s	80%	30s	100%	16s	100%	13s
Exam 7	31	50	F	93.55%	62s	87.09%	59s	80.64%	25s	100%	18s
Exam 8	9	54	F	100%	12s	100%	12s	100%	7s	100%	5s
Exam 9	22	51	F	95.45%	55s	95.45%	55s	77.27%	17s	100%	13s
Exam 10	97	53	F	92.78%	135s	91.75%	132s	77.31%	83s	100%	61s
Exam 11	38	48	F	92.10%	61s	84.21%	57s	78.94%	38s	100%	32s
Exam 12	22	68	F	95.45%	34s	81.81%	33s	90.90%	18s	100%	14s
Exam 13	20	52	F	95%	38s	60%	38s	80%	16s	100%	11s
Exam 14	75	63	M	80%	125s	62.66%	122s	45.33%	78s	100%	60s
Exam 15	22	66	F	95.45%	35s	90.90%	35s	86.36%	20s	100%	14s
Exam 16	50	37	F	88%	99s	82%	98s	76%	48s	100%	37s
Exam 17	18	43	F	88.88%	28s	88.88%	27s	83.33%	17s	100%	13s
Exam 18	55	62	M	98.18%	94s	92.72%	93s	40%	45s	100%	32s
Exam 19	83	70	M	93.97%	141s	90.36%	139s	36.14%	86s	100%	69s
Exam 20	45	64	M	95.55%	72s	86.67%	70s	33.33%	45s	100%	35s
Exam 21	40	56	F	87.5%	69s	80%	68s	95%	43s	100%	34s
Exam 22	44	79	M	84.09%	77s	79.54%	76s	54.54%	43s	100%	33s
Exam 23	37	55	M	86.48%	66s	78.37%	64s	56.75%	44s	100%	35s

In general, when the other contours were detected wrongly, the major cause was the mislabeling of the neighboring *reference objects* such as the lamina. The problems in most cases arise in CT exams where the standard acquisition protocol had not been followed such that one or more unexpected objects (e.g., arms) were present in the images. The presence of such objects affects the symmetry of the CT image.

B. Computational Cost

Our algorithms were implemented on a PC Windows machine, with a processor Pentium III 500-MHz, 512-MB RAM. As presented in Table II, the spinal cord detection is obtained practically in real time. The worst case is when the snakes are used in all the slices. The required amount of interaction is minimal and most often consists of manual correction of the errors in a couple of slices.

In the Anatomic Structures Map, about ten different structures are represented so the query process performed by the plan solver terminates quickly. The most time consuming routine is the Snake algorithm, which optimizes the location of the contour by minimizing the total energy. All other routines are performed in less than a second per slice (0.3–0.5 s) in our current application. The snake algorithm is only applied to slices where the spinal cord is not surrounded by the spinal canal, which boils down to about half of the slices in a typical patient.

VII. DISCUSSION

The major contribution of this article is that a top-down knowledge-based system has been developed for flexible interpretation of CT images. Our system can cope with a large amount of inter-patient variation. Not only is the size and shape of the tumor unique to each patient, also the acquisition parameters of the CT scan vary considerably. The ASM is a means to represent, in a compact form, important fragments of the anatomic knowledge that is used by a radiologist while interpreting the CT images. The plan solver contains the procedural knowledge: how to detect particular anatomical structures one-by-one. The knowledge-based architecture often makes it possible to cope with exceptional conditions which occur in a specialized clinic. Even if our approach fails, it is possible for the radiologist to "take-over" and correct the wrongly placed contours of, e.g., the lamina. Moreover, it is possible to identify the cause(s) of failure because of the transparent knowledge-based architecture.

The major drawback of our approach is the time it takes to formalize the anatomical and procedural knowledge that is needed to implement the ASM and the plan solver. This problem is well-known from the research in expert systems and has been called the *knowledge elicitation bottleneck*. Although the ASM may partly be reused for, e.g., automatic interpretation of MR images of the thorax, reverse engineering would be required to tailor

the plan solver such that is applies the appropriate atomic (segmentation) algorithms. It is clear that a different image modality will, in general, require different low-level operators to find the same anatomic structures. We furthermore wish to add that the rather confined macro-anatomy of the human thorax makes it well-suited for representation in a frame-like hierarchical representation scheme such as the ASM. Representation of, e.g., the human vascular system would be virtually impossible.

A final question addresses the transferability of the developed method to other hospitals. The *DICOM* standard (respected by our system), and the use of HU should work in a hospital with well-calibrated CT scanners. However, we are about to test the system in another hospital in Switzerland. About the feature of the system, probably the most important one is the use of a knowledge-based image processing philosophy, keeping in mind that relying solely on the classical image processing algorithms are insufficient for detecting automatically anatomical structures in CT images. Another important aspect of this approach is the way the 3-D processing is performed, resembling the procedure performed by the radiologist. Finally, this system was tested on images from 23 real patients obtained from a local hospital in Switzerland, and the results were evaluated by specialists in radiology.

VIII. Conclusion

Radiotherapy of malignant tumors located in the vicinity of the is spinal cord requires a very accurate planning to avoid causing unnecessary damage in this vital organ. The spinal cord is a highly radiosensitive organ; even moderate doses of radiation can cause different complications such as paralysis of the patient. In this article, we present a knowledge-based approach to interpretation of CT images. The approach is based on two closely linked knowledge bases: the ASM and the plan solver. The former represents structural (static) knowledge of the macro anatomy in the human thorax. The latter represents the procedural knowledge—the scripts that are used for detection of the different objects of interest. The plan solver combines atomic and composite image processing operators using an inheritance scheme. Which (composite) operators inherit an atomic operator, say a snake algorithm, is derived from the ASM, which contains the structural knowledge.

The knowledge-based approach was implemented on a standard PC. The system was subsequently validated on CT image data from 23 patients who were to undergo radiotherapy. The plan solver was used to locate the following four kinds of objects: the spinal cord, the spinal canal, the lamina, and the body (outer thorax). The highest accuracy was obtained for the body, which was located correctly in all slices among the 23 patients. The spinal canal was located with an accuracy of 92%, the spinal canal with an accuracy of 85% and the lamina with an accuracy of 72%.

The major advantage of our knowledge-based system compared with state-of-the-art low-level solutions lies in its transparency and its flexibility. The system is transparent to the radiologist because parts of his/her medical knowledge is represented in the ASM and the plan solver. Transparency makes it easier to take over from the system in case the identification of the objects fails. Flexibility is required because the scan protocol varies among the patients depending on the location and size of the tumor.

References

[1] H. K. Awwad, *Radiobiological and Physiological Perspectives—The Boundary Zone Between Clinical Radiotherapy and Fundamental Radiobiology and Physiology*. Norwell, MA: Kluwer Academic, 1990.

[2] E. J. Hall, *Radiobiology for the Radiologist*. Baltimore, MD: Williams & Wilkins/Lippincott, 2000.

[3] T. L. Philips and S. A. Leibel, *Textbook of Radiation Oncology*. Philadelphia, PA: Saunders, 1998.

[4] J. Bijhold, K. G. A. Gilhuijs, M. vanHerk, and H. Meertens, "Radiation-field edge-detection in portal images," *Phys. Med. Biol.*, vol. 36, no. 12, pp. 1705–1710, 1991.

[5] K. G. A. Gilhuijs, P. J. H. vandeVen, and M. vanHerk, "Automatic three-dimensional inspection of patient setup in radiation therapy using portal images, simulator images, and computed tomography data," *Med. Phys.*, vol. 23, no. 3, pp. 389–399, 1996.

[6] G. M. te Brake and N. Karssemeijer, "Segmentation of suspicious densities in digital mammograms," *Medical Physics*, vol. 28, no. 2, pp. 259–266, 2001.

[7] B. van Ginneken and B. M. ter Haar Romeny, "Automatic segmentation of lung fields in chest radiographs," *Med. Phys.*, vol. 27, no. 10, pp. 2445–2455, 2000.

[8] N. Archip and P.-J. Erard, "Anatomical structures map—A way of encoding medical knowledge in a task oriented computed tomography image analysis," in *Proc. 2000 Int. Conf. Mathematics and Engineering Techniques in Medicine and Biological Sciences*, Las Vegas, NV, June 2000, pp. 377–383.

[9] O. Wegener, *Whole Body Computed Tomography*, 2nd ed. Oxford, U.K.: Blackwell Scientific, 1992.

[10] M. J. F. Grimnes, "ImageCreek: A knowledge level approach to case-based image interpretation," Ph.D. dissertation, Univ. Trondheim, Trondheim, Norway, 1998.

[11] W. Menhardt *et al.*, "Knowledge based interpretation of cranial MR images," Philips GmbH, Forschungslaboratorium Hamburg, Hamburg, Germay, Tech. Rep., 1990.

[12] M. Sonka, G. Sundaramoorthy, and E. A. Hoffman, "Knowledge-based segmentation of intrathoracic airways from multidimensional high resolution ct images," in *Proc. Medical Imaging 1994: Physiology and Function From Multidimensional Images*, pp. 73–85.

[13] G. L. Vernazza *et al.*, "A knowledge-based system for biomedical image processing and recognition," *IEEE Trans. Circuits Syst.*, vol. 34, pp. 1399–1416, Nov. 1987.

[14] A. Zahalka and A. Fenster, "An automated segmentation method for three-dimensional carotid ultrasound images," *Phys. Med. Biol.*, vol. 46, pp. 1321–1342, 2001.

[15] L. Jennifer, T. E. Boes, C. Weymouth, and R. Meyer, "Multiple organ definition in ct using a Bayesian approach for 3-D model fitting," in *Vision Geometry IV, Proc. SPIE 2573*, 1995, pp. 244–251.

[16] V. P. Kumar and U. B. Desai, "Image interpretation using Bayesian networks," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, pp. 74–77, Jan. 1996.

[17] R. Tombropoulos *et al.*, "A decision aid for diagnosis of liver lesions on MRI," Section on Medical Informatics, Stanford Univ. Sch. Med., Stanford, CA, Tech. Rep., 1994.

[18] F. Kummert *et al.*, "Control and explanation in a signal understanding environment," *Signal Processing*, vol. 32, pp. 111–145, 1993.

[19] E. Rogers, "Via-rad: A blackboard-based system for diagnostic radiology," *Artificial Intell. Med.*, vol. 7, pp. 343–360, 1995.

[20] M. S. Brown *et al.*, "Knowledge-based segmentation of thoracic computed tomography images for assessment of split lung function," *Med. Phys.*, vol. 27, no. 3, pp. 592–598, 2000.

[21] H. Li, R. Deklerck, B. De Cyper, A. Hermanus, E. Nyssen, and J. Cornelius, "Object recognition in brain ct-scans: Knowledge based fusion of data from multiple feature extractors," *IEEE Trans. Med. Imag.*, vol. 14, pp. 343–360, June 1995.

[22] K. Englmeier *et al.*, "Model based image interpretation of spiral ct scans of the abdomen," in *Artificial Intell. Med.*: IOS Press, 1993, pp. 44–51.

[23] L. Gong and C. A. Kulikowski, "Composition of image analysis processes through object-centered hierarchical," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, Oct. 1995.

[24] M. Egmont-Petersen, "Mental models as cognitive entities," in *Proc. Scandinavian Conf. Artificial Intelligence*, B. Mayoh, Ed., 1991, pp. 205–210.

[25] M. Egmont-Petersen, P. C. W. Hogendoorn, R. v. d. Geest, H. A. Vrooman, H. J. v. d. Woude, J. P. Janssen, J. L. Bloem, and J. H. C. Reiber, "Detection of areas with viable remnant tumor in postchemotherapy patients with Ewing's sarcoma by dynamic contrast-enhanced mri using pharmacokinetic modeling," *Magn. Reson. Imag.*, vol. 15, no. 5, pp. 525–535, 2000.

[26] "Prescribing, recording and reporting photon beam therapy," Int. Commission Radiation Units and Measurements, Bethesda, MD, 1993.

[27] J. A. Purdy and G. Starkschall, *A Practical Guide to 3D Planning and Conformal Radiation Therapy*. Sun Prairie, WI: Ed. Advanced Medical Publishing, Inc., 1999.

[28] P. Jackson, *Introduction to Expert Systems*. Reading, MA: Addison-Wesley, 1999.

[29] N. Archip, P. J. Erard, J.-M. Haefliger, and J. Germond, "Lung metastaces detection and visualization on ct images—A knowledge-based method," *J. Visualization Comput. Animation*, vol. 13, no. 1, pp. 65–76, 2002.

[30] R. Adams and L. Bischof, "Seeded region growing," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 16, pp. 641–647, June 1994.

[31] J. R. Parker, *Algorithms for Image Processing and Computer Vision*. New York: Wiley Computer, 1997.

[32] M. Held, "Vroni: An engineering approach to the reliable and efficient computation of Voronoi diagrams of points and line segments," *Computational Geometry*, vol. 18, pp. 95–123, 2001.

[33] D. T. Lee, "Medial axis transformation of a planar shape," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-4, pp. 363–369, Apr. 1982.

[34] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1988.

[35] D. J. Williams and M. Shah, "A fast algorithm for active contours and curvature estimation," *CVGIP: Image Understanding*, vol. 55, no. 1, pp. 14–26, 1992.