Automatic Segmentation of Cervical Vertebrae in X-Ray Images

Xi Xu, Hong-Wei Hao, Xu-Cheng Yin *
School of Computer and Communication Engineering
University of Science and Technology Beijing
Beijing 100083, China
xuxihasan@yahoo.com.cn, hhw@ustb.edu.cn,
xuchengyin@ustb.edu.cn

Ning Liu
Corporate Technology
Siemens Ltd.
Beijing 100102, China
liuning@siemens.com

Shawkat Hasan Shafin
School of Computer Science and Engineering
Beijing University of Aeronautics and Astronautics
Beijing 100083, China
shafin878@yahoo.com

Abstract—**Physiological parameters of vertebrae are important for cervical condition assessment. In order to measure the parameters fast and accurately, automatic segmentation instead of manual key point placement has become an imperative for diagnosing. We propose an applicable automatic segmentation system for medical image of cervical spine. The system includes a series of algorithms: a parallel cascade structure based Haar-like features and the AdaBoost learning algorithm used to detect the location of cervical vertebrae as a initial position of Active Appearance Model (AAM), multi-resolution AAM search applied to improve the speed and accuracy of AAM fit, and combination of global AAM and local AAM used to achieve more effective matching of details of vertebrae. Experiments on the cervical spine databases show a significant increase in speed, robustness and quality of fit compared to previous methods.**

Keywords-cervical vertebra; haar-like feature; parallel cascade structure; multi-resolution AAM search;local AAM

I. INTRODUCTION

Recent surveys have shown that more and more people are suffering from cervical spondylopathy, especially white collar workers and senior citizens. While, analysis of cervical condition is mostly based on medical images as many physiological parameters can only be observed from the images. Thus, how to measure parameters fast and accurately is becoming highly significant.

For a long time, to measure parameters, key points, which mark out several regions of interests, have to be selected manually by medical workers. Obviously, manual point placement is arduous, imprecise and subjective. In addition, with increasing use of medical images, the task is becoming time-consuming and identification of similar cases is getting impossible. If key points are marked by computer automatically, it will be very easier to measure large numbers of cervical parameters and case database can be built to assist in diagnosing.

Segmentation is the process through which the contours of cervical vertebrae are identified and marked for further examination. Medical images may provide noisy and possibly incomplete data, and typically deal with complex and variable structure. Model based methods offer solutions to these difficulties. Widely used segmentation model are Snake, Level set, Active Shape Model and Active Appearance Model (AAM) [1]. By experiments and comparison, AAM is chosen in our research as it is rapid, accurate and robust. It uses a complete model of the image grey level structure, and also models the covariance of that structure. As is well known, AAM is sensitive to the initial position. If the initial position is not proper, AAM will converge to an illogical location as it searches just near the initial position. Many researches have been done to estimate initial position automatically, but most of them have much manual adjustment and suffer the influence of image quality, noise and edges. Furthermore, the global AAM is likely to be over-constrained as it is based on finite training samples and will fail to adequately fit unseen cervical vertebrae which deviate too far from the samples.

To solve these problems, many improvements have been made. We first develop a parallel cascade structure based Haar-like features and the AdaBoost [2] learning algorithm to detect the initial position of AAM. The structure consists of three classifier cascades which are trained respectively using training samples to estimate both image type and initial position. Then, three AAMs are created for each image type to detect feature points. For each AAM, multi-resolution matching which is implemented by a Pyramid model is applied to improve the efficiency and robustness of AAM. Global AAM is combined with local AAMs to achieve more effective matching of details of X-Ray image. Computers may be able to more quickly, accurately, and cost-effectively mark the shapes of the spine by

* Corresponding author: Xu-Cheng Yin.
The work is partly supported by the National Natural Science Foundation of China (61105018, 61175020), the R&D Special Fund for Public Welfare Industry (Meteorology of China (GYHY201106039, GYHY201106047), and the Fundamental Research Funds for the Central Universities (FRF-BR-10-034B).

using the improved segmentation technique. The detailed vertebral shape and appearance parameters of the models obtained in automatic computer determination could provide a quantified form of some of the more subtle aspects of visual or semi-quantitative for expert reading and classifying of cervical spinal disease and conditions. Therefore, this research strives to dramatically decrease the amount of manual intervention required in locating the vertebrae and their key characteristics.

The remaining section of the paper is organized as follows. Section II will describe the method to detect initial position. Section III will detail the segmentation using improved AAMs. Section IV will show a number of experiment results. Finally, Section V contains a conclusion and discussion of our work.

II. DETECTION OF INITIAL POSITION

As our research is based on practical medical diagnosis, first of all, we present some specific descriptions of cervical vertebrae in X-ray images.

In cervical spine examination, main types of X-ray images, which are taken from the same patient, include frontal posture, lateral posture, extension (bending forward) and flexion (bending backward). Here, we focus on the last three types (see Fig. 1). In Fig. 1 the left one is lateral posture, tagged "0"; middle, extension, "1"; right, flexion, "2".

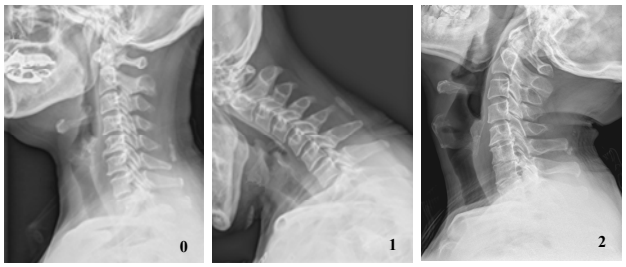

Figure 1. X-ray images of cervical spine.

The cervical spine has seven vertebrae as shown in Fig. 2, labeled from C1 to C7. C1 is a ring that does not have the vertebral body. It is attached to C2. Therefore, C1 is not normally used in most vertebral assessment.

Figure 2. Structure of cervical spine.

The spinous process of a vertebra locates behind the corresponding vertebra body. Accurate contours of vertebra bodies and spinous processes are highly significant for cervical spine examination. Compared with vertebral bodies, spinous processes are various, irregularly shaped, and often fuzzy because of low image quality; in addition, C7 is usually has an over-bright "washed-out" appearance. However, our purpose is to fix a key point set of vertebral bodies and spinous processes for each image. Thus, the segmentation is more challenging for us.

A. Weak Classifier based on Haar-like Features

We use the value of simple features to recognize initial position. Haar-like features, reminiscent of Haar basis functions, are rectangle features. They can indicate specific characteristics in an image and have been used for human face detection by Papageorgiou et al. [3].

Each Haar-like feature consists of several jointed black and white rectangles as shown in Fig. 3. The value of a Haar-like feature is the difference between the sums of the pixels gray level values within the black and white regions. Rectangle features can be computed very rapidly using the integral image, which is also called a summed area table [4].

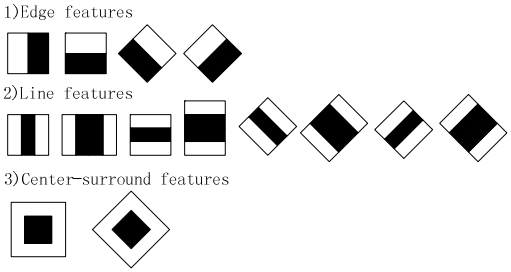

Figure 3. A set of extended Haar-like features.

The integral image at location x, y contains the sum of the pixels above and to the left if x, y, inclusive:

$$ii(x,y)=\sum_{x'\le x,y'\le y} i(x',y') \quad , \tag{1}$$

where *ii(x, y)* is the integral image and *i(x, y)* is the original image (see Fig. 4(a)). The integral image can be computed in one-pass over the image using the following pair of recurrences [5]:

$$s(x,y)=s(x,y-1)+i(x,y)\,, \tag{2}$$

$$ii(x,y)=ii(x-1,y)+s(x,y)\,. \tag{3}$$

The sum of pixels within "D" can be computed by (see Fig. 4(b)):

$$\mathrm{Re}ctsum(D)=ii_4+ii_1-ii_2-ii_3\,. \tag{4}$$

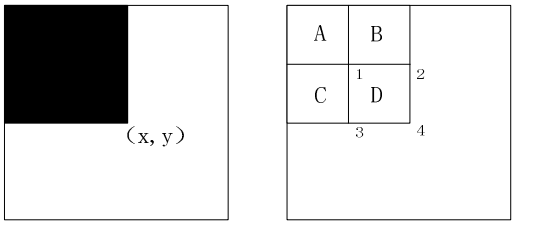

(a) Integral value of (x,y) (b)Compute integral value of D

Figure 4. Computation of Haar-like features.

Therefore, the Haar-like features have the advantage of very fast computation, because it depends only on the sum of pixels within a rectangle instead of every pixel value. Haar-like features capture the intensity gradient at different locations, spatial frequencies and directions by changing the size, shape, position and arrangement of rectangular regions exhaustively according to the base resolution of the detector. Given that the base resolution of the detector is 24×24, the exhaustive set of features is quite over 180,000 [6]. The features are sensitive to the presence of edges, bars, and other simple image structure. This allows them to be used to detect objects of various sizes.

To detect the cervical spine, the X-ray image is scanned by a sub-window x containing a Haar-like feature. For each feature f, the weak learner determines the optimal threshold θ classification function that misclassifies the minimum examples. A weak classifier *h(x, f, g, θ)* is designed to select the single rectangle feature which best separates the positive and negative examples:

$$h(x, f, g, \theta) = \begin{cases} 1, pf(x) < p\theta \\ 0, otherwise \end{cases}, \quad (5)$$

where p indicates the direction of the inequality sign.

B. AdaBoost learning Algorithm

Recall that in each sub-image, the total number of Harr-like features is far larger than the number of pixels. Although calculating a feature is extremely fast and efficient, calculating all features is time-consuming and expensive. Moreover, only a very small number of the features can be combined to form an effective classifier. Thus, a variant of AdaBoost is used in our research to select these features and train the classifiers [2]. AdaBoost is a powerful machine learning algorithm that can learn a strong classifier based on a large set of weak classifiers by re-weighting the training samples. The learning algorithm is as follows:

- Step 1: Initialize the weights w_t over training examples with a uniform distribution. The weights tell the learning algorithm the importance of the example.
- Step 2: Choose the classifier $h_t(x)$ with the lowest error ε_t for each feature by using the weak learning algorithm *h(x, f, g, θ)*.
- Step 3: Deduce the weights on the training examples that were classified correctly to $\frac{w_t \varepsilon_t}{1-\varepsilon_t}$.
- Step 4: Repeat Step 2 until T weak classifiers are selected.
- Finally, make a linear combination of the weak classifiers obtained at all iterations, $h_1(x)$, …, $h_T(x)$ according to (6):

$$C(x) = \begin{cases} 1, \sum_{t=1}^{T} a_t h_t(x) \geq \frac{1}{2} \sum_{t=1}^{T} a_t \\ 0, otherwise \end{cases} \quad (6)$$

where $a_t = \log \frac{1-\varepsilon_t}{\varepsilon_t} \approx -\log \varepsilon_t$ (when ε_t is quite small). *C(x)* is a strong classifier. A more comprehensive description of the formula is give by Freund et al [2].

C. Parallel Cascade Structure

A cascade of classifiers achieves increased detection performance while radically reducing computation time [6]. It is degenerated decision tree [7]. Each stage, which actually is a strong classifier, is trained using AdaBoost learning algorithm. A positive result from the first stage triggers the evaluation of a second stage, which has also been adjusted to achieve very high detection rates. A positive result from the second stage triggers a third stage, and so on. A negative outcome at any point leads to the immediate rejection of the sub-image. As the stage goes further, the number of weak classifiers, which are combined to form the strong classifier to achieve the desired detection rate, increases. Therefore, more computation is required. However, after several stages of processing, the number of negative sub-images has been reduced radically (see Fig. 5). "true" means a positive result; "false" means a negative outcome.

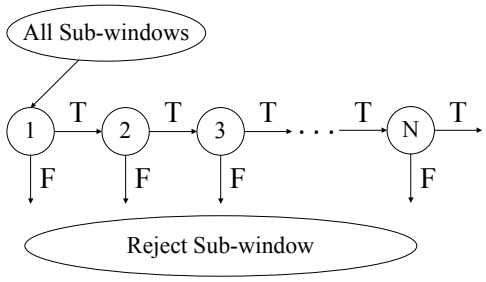

Figure 5. A cascade of classifiers with N stages. T denotes "true"; F denotes "false".

For each posture of cervical spine, a classifier cascade is trained based on the AdaBoost learning algorithm to achieve the required accuracy. We construct a parallel structure that includes all three classifier cascades (see Fig. 6).

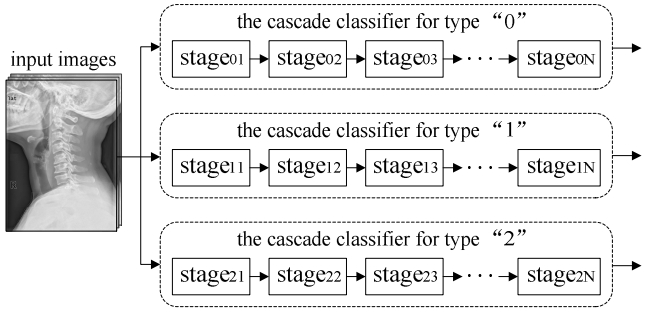

Figure 6. A parallel structure of three cascades.

The input image is detected by all classifier cascades respectively at the same time. The image type depends which cascade classifier has an output. In the output image, the initial position is marked with a rectangular region, containing all vertebrae from C2 to C7.

III. IMPROVEMENTS OF ACTIVE APPEARANCE MODEL

After a rough initial position, which determines where AAMs should be placed in an image, is estimated, we now turn to details of AAMs and some improvements in practice.

A. Active Appearance Model

Active Appearance Model (AAM) is a parametric deformable model. It is generated by combining a model of shape variation with a model of texture variation [8].

1) Shape model

To build a shape model, we require a training set of cervical spine image marked with control points defining the main features (see Fig. 7). The process is critical and it determines the efficiency and accuracy of AAM matching.

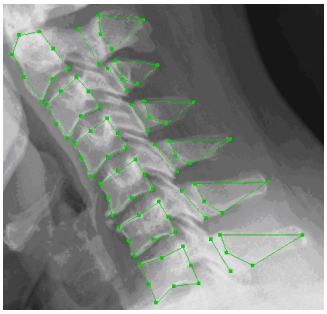

Figure 7. Control points of cervical vertebrae.

The control points are labeled on the contour of the cervical vertebrae and some special lines. A vertebral body has eight control points: four at the corner and the other four at the middle of the outline; each spinous process gets four points: three at front and one at tail. Our experiments have proved that if two control points is added between the vertebral body and its corresponding spinous process, AAM usually converges to a more accurate location of spinous process. The points are along the bright line between the vertebral body and the spinous process. Thus, each vertebral contour uses 84 points in all. If a point is described by 2-dimensional coordinate, a set of control points can be represented as a vector with length of 168:

$$X = (x_1, y_1, x_2, y_2, ..., x_{84}, y_{84})^T. \quad (7)$$

All shape vectors of the training set are align into the mean shape vector and a statistical shape model is built by applying Principal Component Analysis (PCA) to aligned shape vectors [9]. Any example can be approximated using:

$$X = \bar{X} + P_s b_s, \quad (8)$$

where $\bar{X}$ is the mean shape vector, P_s is a matrix consisting of the eigenvectors of the shape covariance matrix, and b_s is a vector of shape parameters. Various shapes can be generated by changing the elements of b_s.

2) Texture model

To form a complete model, the texture also needs to be considered. We warp each training image using a triangulation algorithm to match its control points with the mean shape, obtaining a "shape-free patch".

The Delaunay triangulation realized by Bowyer-Watson algorithm is applied in our research to perform image warping (see [10] for details). It connects the set of points using a triangular mesh such that each triangle satisfies the Delaunay property [11]. Fig. 8 shows the result of Delaunay triangulation.

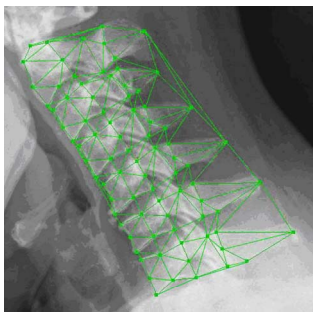

Figure 8. Result of Delaunay triangulation.

We then sample the intensity information from the shape-normalized image over the region covered by the mean shape. The resulting samples are normalized by applying linear transformations to minimize the influence of lighting and exposal [8]. Finally, we apply PCA to the normalized data to build a texture model:

$$G = \bar{G} + P_g b_g, \quad (9)$$

where $\bar{G}$ is the mean normalized grey-level vector, P_g is a matrix consisting of the eigenvectors of the texture covariance matrix, and b_g is a set of grey-level parameters.

3) Appearance model

For each example we concatenate the shape and texture vectors and obtain a new vector:

$$B = \begin{pmatrix} W_s b_s \\ b_g \end{pmatrix} = \begin{pmatrix} W_s P_s^T (X - \bar{X}) \\ P_g^T (G - \bar{G}) \end{pmatrix}, \quad (10)$$

where W_s is a diagonal matrix of weights for each shape parameter, allowing for the difference in units between the shape and texture models [1]. After a further PCA is applied on b to remove correlation between shape and texture model parameters, we obtain a combined appearance model:

$$B = \begin{pmatrix} Q_s \\ Q_g \end{pmatrix} c = Qc, \quad (11)$$

where Q is a matrix consisting of the eigenvectors of b and c is a vector of appearance parameter, controlling both the shape and texture of the model according to:

$$X = \bar{X} + P_s W_s^{-1} Q_s c = \bar{X} + R_s c \quad (12)$$

and

$$G = \bar{G} + P_g Q_g c = \bar{G} + R_g c, \quad (13)$$

where $\bar{X}$ is the mean shape, $\bar{G}$ is the mean texture in a mean shaped patch, and R_s, R_g are matrices describing the models of variation derived from the training set [1].

4) AAM search

When a full appearance model as described above is built and a rough initial position of model is estimated, various synthetic examples can be generated by adjusting the model parameters. AAM search is to find a synthetic example which matches a given novel image as closely as possible, namely minimizing the current difference between model and image (measured in the normalized texture frame). That is minimize

$$\Delta = \|\delta_G\|^2 = \|G_n - G_m\|^2, \tag{14}$$

where G_n is obtained by sampling the pixels in this region of the image and projecting into the texture model frame. G_m is the current model texture and calculated by (13). A more comprehensive description is give by Coote et al [8].

B. Multi-resolution AAM Search

In order to improve the speed and accuracy of image interpretation, we use a multi-resolution matching based on Gaussian image pyramid. It is a coarse-to-fine strategy. AAM is applied first to a coarse, low resolution version of the image, and then refined on higher resolution versions. A multi-resolution pyramid like Fig. 9 is constructed for each image. It has several layers. The image at Level 0, basic layer, is the original image; the image at Level 1 has half the number of pixels of that at Level 0, which is generated using Gaussian smoothing and sub-sampling from the original; the images of other levels are obtained in the same way.

For each image, we generate a set of AAMs, one for each level of the pyramid we wish to use. That means we built a separate AAM for each layer. As the pixels in Level L are 2^L times the size of that in Level 0(see Fig. 9), the points on higher level will have grander adjustment during one iteration and the model can be located near target object quickly, while the lower level allows only small movement for more accurate location. A three-level multi-resolution pyramid is applied in our research. Fig. 10 and Fig. 11 show the pyramids of the shape and the texture models of cervical spine respectively.

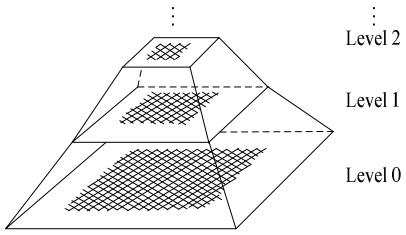

Figure 9. Pyramid model of image.

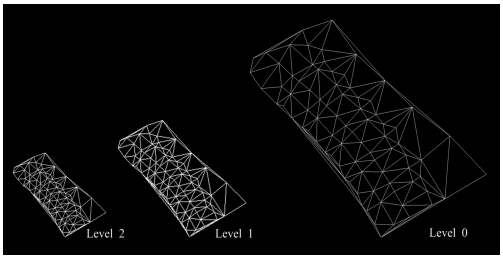

Figure 10. Pyramid of shape model.

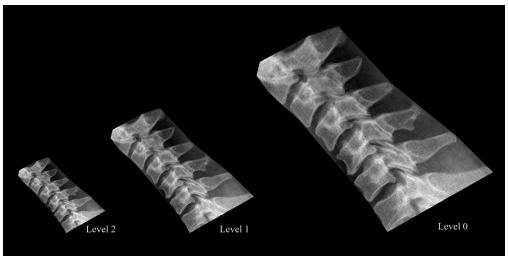

Figure 11. Pyramid of texture model.

Local search is carried out top-down. Limits are set on each level to decide whether stop iterations of the current level and move to the next lower level in the pyramid. The multi-resolution AAM search algorithm is described as follows:

- Step 1: Start at the highest level of the pyramid.
- Step 2: Search at the current level. Run a number of iterations of AAM using the models trained at the current level until the model has converged or the limit number of iterations has been reached.
- Step 3: If the current level is not Level 0, move to the next level down in the pyramid, project the current model into the updated level and repeat from Step 2; otherwise, finish the algorithm.

Our experiments have proved that the algorithm is faster and can converge to a more accurate location from further away than searching at a single resolution.

C. Combination of Global AAM and Local AAM

To achieve more effective fit of details, we build both a global model of the whole cervical spine, and a sequence of local AAMs composed of overlapping vertebral triplets. An initial global model provides further information for the local model starting solutions. Each vertebra is matched using the triplet model in which it is central, and the two neighboring vertebrae provide helpful constraints and linkage to the other local models [12]. Four triplet models are respectively built covering the cervical spine from C2 down to C7, namely C2/C3/C4, C3/C4/C5, C4/C5/C6, and C5/C6/C7. Fig.12 demonstrates a triplet model. There are 42 feature points in each triplet model. It has been proved that triplets of vertebrae are the optimal structure to model [13]. If the modeled structure is too larger, the local model has the same problem as global AAM; whereas a smaller structure is too unconstrained and causes distortion in shape.

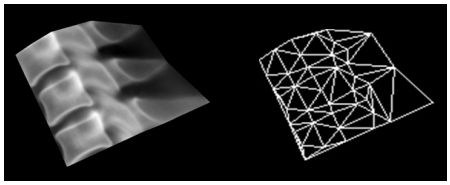

Figure 12. A triplet model (C3/C4/C5).

The combination of global AAM and local AAM consists of the following. First, match the global shape model to the initial position as a starting solution. Then, perform local AAM

fit in order of C2/C3/C4, C3/C4/C5, C4/C5/C6, and C5/C6/C7. The concrete steps for each triplet model are as follows:

- Step 1: Project points from the global shape vector into the local model. Perform a fit of the local shape model to these projected point positions.
- Step 2: Iterate a multi-resolution AAM fit of local model to the image (see Fig. 13) and project back the outcome to the local shape vector.
- Step 3: Project back the updated points of both the middle and the bottom vertebrae into the global vector. The points of the top vertebra are not projected back.
- Step 4: Perform a fit of the global shape model to the updated global vector. Copy out from the global model any as yet unfitted points into the global point vector.

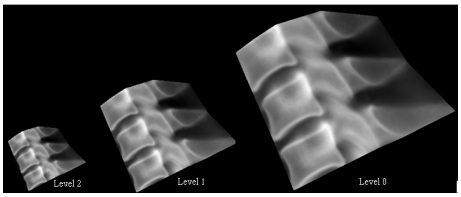

Figure 13. Three-level pyramid of local Active Appearance model.

At each triplet model fit iteration, the current solution is projected from the global model into the local model, and then the local model fit is projected back into the global model [12]. Current model determines the points of its middle vertebra and adjusts some of the initial points of the next local model. The local AAM can avoid misplacement of vertebrae and reduce the influence of the lighter area of shoulder.

IV. EXPERIMENTS

Fig. 14 shows the workflow of our automatic segmentation algorithm for cervical vertebrae image.

We have collected 10024 X-ray images of cervical vertebrae from various sources as our database. 9357 images used as training samples were from well-known database, NHANES II [14], and the rest used as test samples were supplied by a first class hospital in Shanghai, China. All images were classified into three groups according to spinal postures.

To train the classifiers, we need both positive and negative training samples. A positive image contains an instance of cervical spine, but the negative does not. Moreover, the positive images should be Close-ups of cervical vertebrae and have great variety. Each classifier training set had 300 positive samples and 400 negative samples. The positive are selected form NHANES II, and scaled and aligned to a base resolution of 20×20 pixels, while the negative are collected by cropping images containing no cervical spine Fig. 15 shows a group of positive training samples. The possible output images of initial position detector are demonstrated in Fig. 16.

There were 60 samples in each AAM training set. Moreover, training sets used in multi-resolution AAM search and local AAM were built in different way according to

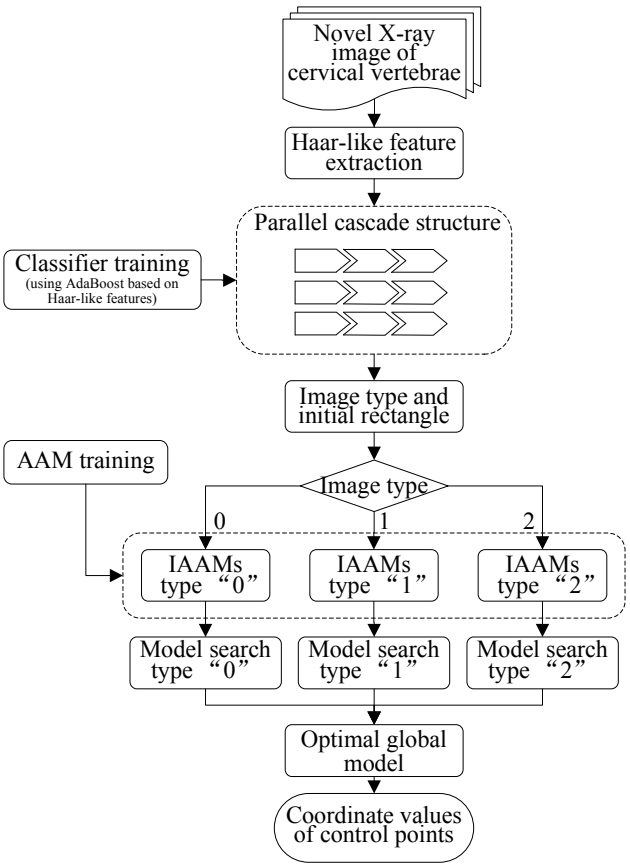

Figure 14. Work flow of our algorithm(IAAMs means improved AAMs).

Figure 15. Some positive training samples of extension.

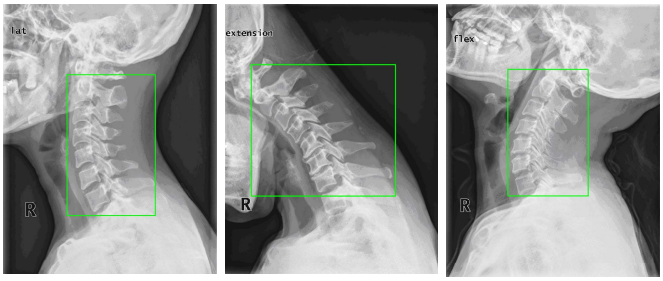

Figure 16. Possible output images for all types.

algorithms. We applied AAM training algorithm to the improved cervical vertebrae models described in section III.

Finally, there were three improved AAMs corresponding to three types of postures.

Fig. 17 compares the search effects of basic AAMs ((a) and (c)) with that of our improved AAMs ((b) and (d)). Although both types of AAMs converge to approximate contours of cervical vertebrae, the improved AAMs obviously fit the details more accurately than traditional AAMs, especially for C6 and C7. That means more precise physiological parameters of vertebrae are obtained to help doctors evaluate the cervical condition.

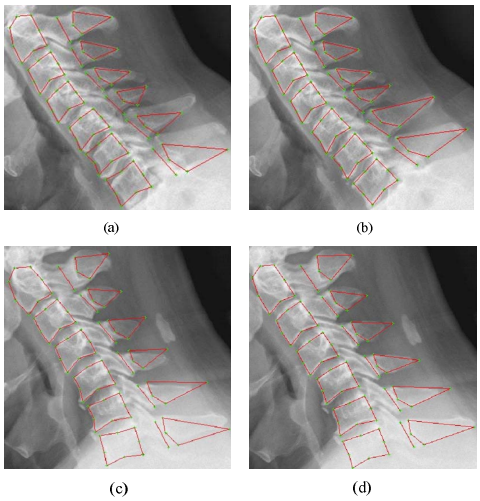

Figure 17. Comparision of AAM search effects.

100 cervical vertebrae images were selected from test samples randomly. Segmentation experiment includes three parts: traditional AAMs (Part I), our improved AAMs (Part II), and our improved AAMs added detection of initial position (Part III). For each sample, we compared the coordinate values of control points obtained from segmentation results with that of corresponding points marked manually by professional doctors, and calculated error by

$$E = \frac{1}{n}\sum_{i=1}^{n}\sqrt{(x_i - x_i^{'})^2 + (y_i - y_i^{'})^2}\,, \qquad (15)$$

where (x_i, y_i) is the coordinate value of control point from the segmentation results, $(x^{'}, y^{'})$ is the coordinate value of corresponding manual marked control point, and n is the number of control point.

The statistical result of experiment is shown in Table I and Fig. 18. The performance of our improved AAMs was obviously better than traditional AAMs, 22 pixels less in average error. When added automatic detection of initial position, the segmentation process no longer needed artificial support. Thus, the speed of segmentation was greatly increased. However, the accuracy rate was slightly declined as 19 samples had large errors (the others' errors were less than 8 pixels). By analyzing the segmentation process, there were two cases resulting in large errors. One was that the rectangular obtained from automatic detection was much bigger than the contour of cervical spine, and the other was C1 or a part of C1 was mistakenly framed in the rectangular.

TABLE I. ERROR STATISTICS

Part of Experiment	Error (pixel)		
	Minimum	Average	Maximum
Part I	8.652	27.597	52.336
Part II	3.682	4.792	6.007
Part III	4.035	19.018	146.925

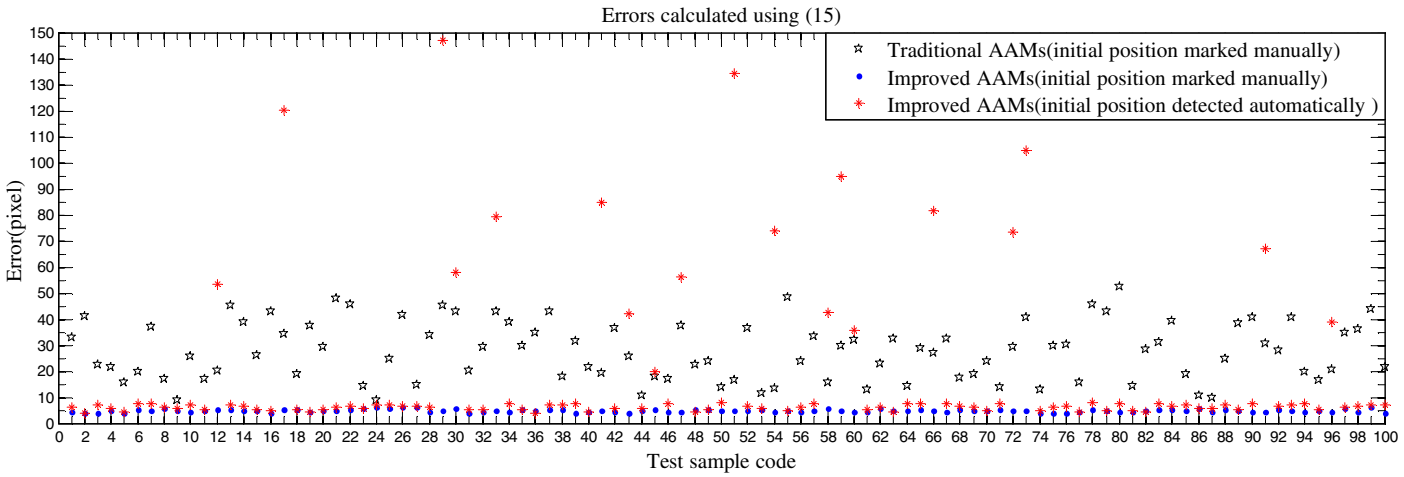

Figure 18. Errors comparison.

V. CONCLUSION AND DISCUSSION

The experiments have proved that the system proposed in our research, including a parallel cascade structure based Haar-like features and the AdaBoost learning algorithm, multi-resolution AAM search, and combination of local AAM and global AAM, have a better performance than traditional AAM both in speed and accuracy of cervical vertebra segmentation. It also indicates that the initial position of AAM is a main factor impacting segmentation effect. Therefore, how to make a good compromise between the automation and precision has become the crux of cervical vertebra segmentation. More research has

to be done to improve the accuracy of the rectangular obtained from automatic detection of initial position.

Achievement of our research on automatic segmentation of cervical spine has already been used in an assistant diagnosis system. The system provides physiological parameters (see Fig. 19 (a), (b) and (c))) and diagnosis advice for doctors (see Fig. 19 (d)).

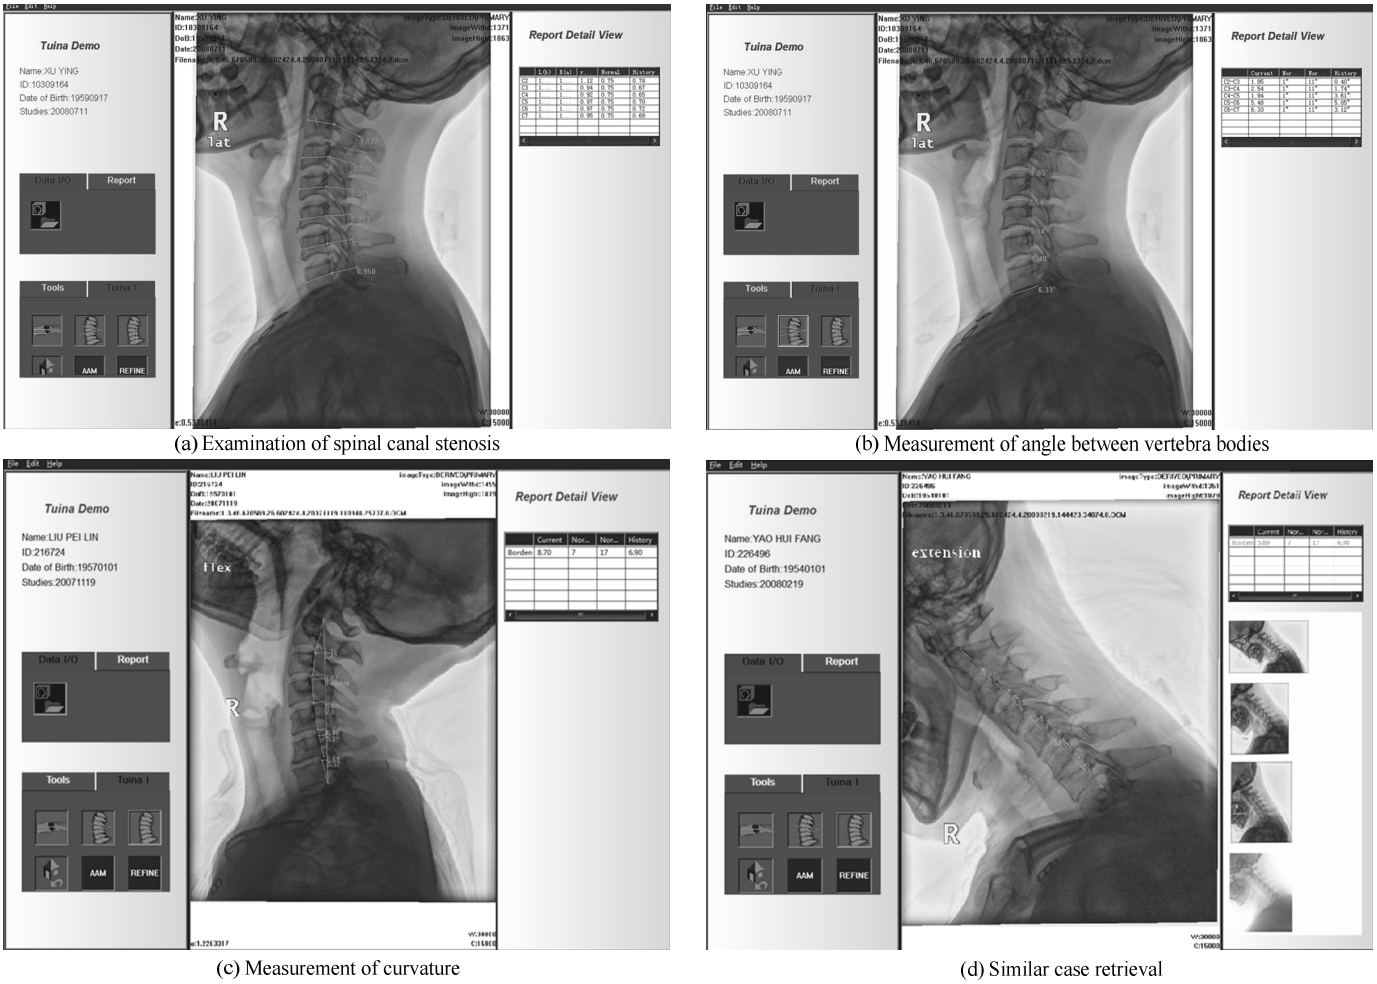

(a) Examination of spinal canal stenosis

(b) Measurement of angle between vertebra bodies

(c) Measurement of curvature

(d) Similar case retrieval

Figure 19. Interfaces of the assistant system.

REFERENCES

[1] T. F. Cootes, and C. J. Taylor, "Statistical models of appearance for medical image analysis and computer vision," in Proc. SPIE Medical Imaging, vol. 3, pp. 138–147, 2001.

[2] Y. Freund, and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," JCSS, vol. 55, pp. 119–139, 1997.

[3] C. P. Papageorgiou, M. Oren, and T. Poggio, "A general framework for object detection," in sixth international Conf. on Computer Vision, pp. 555–562, Jan 1998.

[4] F. Crow. "Summed-area tables for texture mapping," in Proceedings of SIGGRAPH, vol. 18, pp. 207–212, 1984.

[5] P. Viola, and M. Jones, "Robust real-time object detection," in 2nd international workshop on statistical and computational theories, July 2001.

[6] P. Viola, and M. Jones, "Rapid object detection using boosted cascade of simple features," in IEEE computer society Conf. on CVPR, 2001.

[7] J. Quinlan, "Induction of decision trees," Machine Learning, vol. 1, pp. 81–106, 1986.

[8] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," IEEE transactions on PAMI, vol. 23, pp. 681–685, June 2001.

[9] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape model-their training and application," Computer Vision and Image Understanding, vol. 61, pp. 38-59, Jan. 1995.

[10] M. B. Stegmann, Active appearance models: theory, extensions, and cases, 2nd ed., Technical University of Denmark, LYNGBY, 2000.

[11] D. F.Watson, "Computing the n-dimensional Delaunay tessellation with applications to Voronoi polytopes," the computer journal, vol. 24, pp. 167–172, 1981.

[12] M. G. Roberts, T. F. Cootes, and J. E. Adams, "Linking sequences of active appearance sub-models via constraints: an application in automated vertebral morphometry," in 14th British Machine Vision Conference, pp. 349–358, 2003.

[13] M. G. Roberts, T. F. Cootes, and J. E. Adams, "Vertebral shape: automatic measurement with dynamically sequenced active appearance models," in 8th MICCAI, vol. 2, pp. 733–740, 2005.

[14] NHANES II Data, http://www.nber.org/data/national-health-and-nutrition-examination-survey-ii.html, March 2007.