Segmentation and image navigation in digitized spine x-rays

L. Rodney Long[*], George R. Thoma

National Library of Medicine, Bethesda, MD 20894

ABSTRACT

The National Library of Medicine has archived a collection of 17,000 digitized x-rays of the cervical and lumbar spines. Extensive health information has been collected on the subjects of these x-rays, but no information has been derived from the image contents themselves. We are researching algorithms to segment anatomy in these images and to derive from the segmented data measurements useful for indexing this image set for characteristics important to researchers in rheumatology, bone morphometry, and related areas. Active Shape Modeling is currently being investigated for use in location and boundary definition for the vertebrae in these images.

Keywords: x-ray, image processing, Active Shape Modeling, statistical shape analysis, segmentation, NHANES, NLM

1. INTRODUCTION

> The ultimate goal of machine vision is image understanding—the ability not only to recover image
> structure but also to know what it represents.[1]

This paper updates work toward the computer-assisted indexing of a collection of 17,000 digitized x-ray images of the cervical and lumbar spine by content of the images. These images were collected as part of the second National Health and Nutrition Examination Survey (NHANES II) and are archived at the National Library of Medicine. Extensive information has been collected describing the physical and demographic characteristics, as well as the health histories, of the subjects of these x-rays, but no information has been derived from the x-ray images themselves. The fundamental information that we seek from the images is the basic geometric structure of the individual vertebrae and the intervertebral geometry; from this data, secondary information related to rheumatological features of interest, pathology and injury may potentially be derived, with expert assistance. Our goal is to eventually index this collection with the basic quantitative geometric data, and also with high-level biomedical feature interpretations of interest to spine researchers. Examples of high-level features of interest to researchers in rheumatology include the presence and degree of anterior osteophytes, disc space narrowing, and subluxation in the cervical spine images. To achieve this eventual goal of indexing the contents of these images with computer assistance, fundamental problems of obtaining basic geometric orientation within the image contents, reliably identifying and navigating among image landmarks, and segmentation of the vertebrae must be addressed. In previous work we have described the results of experiments to obtain basic anatomy-referenced orientation within the cervical spine images; in this paper we review our approach to obtaining this anatomy-referenced information, discuss an algorithm (Active Shape Modeling) that may use this information to initialize searches for the boundaries of individual vertebrae of the spine, and present initial results of testing this algorithm on cervical spine images. Successful use of Active Shape Modeling for spine images has been reported by Smyth[2] for the lumbar spine for dual x-ray absorptiometry (DXA) scans, but we are not aware of successful implementations for digitized x-ray film images of the spine, although related work using active contours for digitized spine x-rays has been published (for example, see Gardner[3]).

2. APPROACH

The conceptual framework we are following is shown in Fig. 1. We seek an iterative method that, beginning with an initial estimate of vertebral boundaries, will gradually refine these estimates until they satisfy pre-set convergence criteria. The specific method that we are evaluating for boundary estimation is Active Shape Modeling (ASM). For the ASM method to be effective, that is, to converge to reasonable vertebral boundaries, we expect that good initial (a priori) values for boundary shape, position, orientation, and scale will be required. A priori boundary shape is obtainable from the statistics of sampled

---

[*] Email: long@nlm.nih.gov

vertebrae (after transformation to "model space", where the original position, orientation, and scale have been removed). Of the position, orientation, and scale variables, it is vertebral position relative to the image frame that has the most obvious variation among the images. This appears to be largely because of the inconsistent placement of the subjects relative to the image frame; that is, the anatomy of interest shifts substantially, relative to the image frame, from image to image. (Other contributors to the vertebral location variance include variations in individual anatomy and in the angle of bending of the neck, but these factors appear to be second order.) This large variance in vertebral location relative to the image frame complicates the use of sample image data in obtaining good a priori information for locating a particular vertebra in a given image. A method that we have experimented with to overcome this problem is to put a frame of reference within the anatomy itself, with the expectation of removing much of the variation due to the shifting of subjects among images. Specifically, we are seeking to develop a reliable method of fixing an Anatomy Based Coordinate System (ABCS) in the cervical spine anatomy, with one axis fixed approximately tangentially to the skull/jaw posterior and the other (orthogonal) axis intersecting at the approximate mid-point of the spine. Presumably, samples of vertebral locations measured relative to this system would have less variation than samples measured relative to the image frame, and the statistics of these locations may provide useful a priori positions for the ASM boundary location. Results of work we have done on the ABCS algorithm are given in the references[4-5] and are summarized in the following section.

3.   ANATOMY-BASED COORDINATE SYSTEM

The concept of creating reference frames in the form of lines or axes in lateral views of human anatomy is not new. Taveras[6], for example, lists six "principal lines" of the skull in lateral view. Of these, one pair, Reid's base line (anthropological basal line) and the auricular line, form an orthogonal axis system. The use of computer algorithms to locate such reference frames within the anatomy presented in a given image does not appear to be common, however. Although we would prefer to use a reference frame from the anatomy literature, such as the one mentioned above, we have to take into account the characteristics of our image data. In many of our images, we do not have the anatomy present to find Reid's base line, for example, which "joins the inferior edge of the orbit with the superior edge of the external auditory meatus"[6]. The auditory anatomy required for this base line is very faint in many images, and the orbit anatomy, not even present in many. Our alternative has been to define a new coordinate system referenced to anatomy that (1) is present in most or all of the images and (2) appears sufficiently prominent to allow reliable location of the system. The anatomical features that we are relying on for this purpose are the relatively bright skull and jaw posteriors and the brightest point of the mid-spine: conceptually, we place a U axis approximately tangent to protuberances on the back of the skull (nominally, in the vicinity of the inion, the external occipital protuberance, or along the skull boundary between that point and the spine) and on the jaw (nominally, in the vicinity of the mandibular angle, or along the jaw boundary between this region and the chin). In our approach, the precise locations of the tangent points on the skull back and on the jaw that define this axis are not critical, and in fact will vary with individual anatomy; our interest is in obtaining a line which will reliably intersect the upper spine vertebrae at a vertebral location that is statistically predictable. The second axis V in the ABCS system is defined to be orthogonal to the U axis and intersecting U at the point of greatest grayscale value within the spine. As stated above, the utility of the ABCS system, if successful, is that it will give us an orientation within the x-ray, which is fixed in the anatomy itself. This system, presumably, can be a basic reference frame for discovering image subregions containing features of biomedical interest, such as vertebral locations, vertebral boundaries, and geometric measures such as vertebral heights and interveterbral spacings.

The anatomy-based coordinate system described above is "located" using an algorithm with these basic steps: (1) Apply grayscale thresholding to the image using a threshold that preserves the essential shape of the skull base in the jaw region and in the region of the skull posterior; (2) in the thresholded image, locate the skull base region at a coarse level: more precisely, locate the jaw region boundary ("below" the spine), and locate the skull posterior region boundary ("above" the spine); in each of these region boundaries, compute slopes of lines tangent to the boundary; this yields multiple tangent slopes for the jaw region boundary, and multiple tangent slopes for the skull posterior region boundary; (3) find the pair of points, one chosen from jaw region boundary, one chosen from the skull posterior region boundary, that have the closest agreement in slope of tangent line; (4) take these points as defining the coordinate system axis U along the skull base; (5) determine a subinterval R of the U axis which contains the intersection of U with the spine; (we define this region R as the subinterval of the U axis bounded by certain characteristic grayscale patterns--two dark areas marking the spine boundaries, separated by a continuous bright area corresponding to the spine itself); and (6) search within the U axis/spine intersection region R for the point O of maximum brightness; this point defines the origin of the U/V coordinate system and the intersection of the orthogonal axis V with the U axis. Steps in calculating the U/V axes are illustrated in Fig. 2. Figure 2A shows a grayscale cervical spine x-ray image. Figure 2B is the same image, with the boundaries for a grayscale threshold superimposed on the original image. It is the boundaries of this thresholded image that are used in the computation of the U axis. Figure 2C

shows the U axis that has been computed. Finally, Fig. 2D shows the V axis added orthogonally to the U axis, and intersecting the U axis at the point of greatest grayscale value in the line segment determined by the U axis/spine intersection.

To date the ABCS algorithm has performed with mixed results, with subjective evaluation of good axes location less than 50% of the trial cases. As remarked above, the current ABCS algorithm searches the contour of an image which is obtained by grayscale thresholding of the original image: many of the failure cases can be attributed to poor performance of the thresholding, resulting in contours that do not accurately follow the anatomical boundaries in the original image. To address this, we have implemented an adaptive grayscale thresholding technique published by Chow and Kaneko[7] and are currently evaluating its effect on algorithm performance. We have discussed our implementation of the Chow-Kaneko algorithm in previously published work[5].

4. ACTIVE SHAPE MODELING

The purpose of Active Shape Modeling (ASM) is to identify the best match to an object model from possible objects contained within a given image. ASM is an iterative algorithm; at each step the object model is allowed to change its position, scale, and orientation, and to modify its shape independently of these position/scale/orientation changes. The model is statistical in nature and is derived from a sample of instances of the object represented by the model. If M instances are used, with N landmark points, which define the object boundary, collected for each instance, then the data used to build the model is represented by the 2NxM matrix

$$A = \begin{bmatrix} \overline{v}_1 & \overline{v}_2 & ... & \overline{v}_M \end{bmatrix}$$

where the $\overline{v}_i$ are 2N-vectors that list the (x,y) coordinates of the landmarks for instance i in the form

$$\overline{v}_i = \begin{bmatrix} x_1 & x_2 & ... & x_N & y_1 & y_2 & ... & y_N \end{bmatrix}^T.$$

The mean $\overline{\mu}$ and covariance C for the sample data A are computed by

$$\overline{\mu} = \frac{1}{M} \sum_{i=1}^{M} \overline{v}_i \text{ and } C = \frac{1}{M-1} \sum_{i=1}^{M} (\overline{v}_i - \overline{\mu})(\overline{v}_i - \overline{\mu})^T \text{ , respectively.}$$

From $\overline{\mu}$ and C we get the basic information about the shape and variability of the model. $\overline{\mu}$ directly gives the expected (a priori) shape we will use to initialize the ASM search. From C we derive the ways in which we expect the object may vary (independently of variations in position, scale, and orientation), using principal component analysis.

Let

$$\Phi = \begin{bmatrix} \overline{e}_1 & \overline{e}_2 & ... \overline{e}_M \end{bmatrix}$$

be the 2NxM matrix having M eigenvectors of C as columns. (Each $\overline{e}_i$ satisfies $C\overline{e}_i = \lambda_i \overline{e}_i$ for some scalar λ_i; then $(\lambda_i, \overline{e}_i)$ is an eigenvalue/eigenvector pair for C.) The vectors $\overline{e}_i$ are the principal components of C and provide important information about the dispersion of the sample points A about their mean vector $\overline{\mu}$. Specifically, any of the sample vectors may be generated by a linear combination of the $\overline{e}_i$, that is, for each j, (j=1,...,M),

$$\overline{v}_j = \overline{\mu} + \sum_{i=1}^{M} b_i \overline{e}_i$$

for appropriately chosen scalars b_i. Thus, any shape in the sample data A may be thought of as being composed of the sum of the mean shape $\overline{\mu}$ and a linear combination of the $\overline{e}_i$: the $\overline{e}_i$ form a set of basis vectors from which all of the shapes represented by the columns of A may be generated. If the sample vectors in A have been selected to be a good representation for the totality of objects to be searched, then the $\overline{e}_i$ may be expected to approximately generate these (new) shapes as well

as those from which C was derived. In the context of ASM, the $\bar{e_i}$ are called the modes of variation for the data A. It is these modes that control the variability from the mean shape that is allowed in ASM. Further, one usually seeks a reduction in the data to be manipulated by using only a subset of these modes. From statistical theory, it is known[8] that the total variance V in A is given by the sum of the variances along each of the eigenvectors $\bar{e_i}$, and that the proportion of the total variance represented by a particular $\bar{e_i}$ is given by its corresponding eigenvalue λ_i divided by V; this gives rise to the following approach: (1) arrange the eigenvectors in order of decreasing eigenvalue; (2) select a variance threshold T (0<T<1); (3) select the first P eigenvectors so that the sum of their corresponding eigenvalues divided by V is at least T. If T is selected to be 0.95, for example, then our P-mode model would represent at least 95% of the variance of the sample points A. The ASM algorithm exploits the principal component approach above. We collect landmark data from M samples to create the matrix A, compute the sample mean $\bar{\mu}$ and covariance C, compute the modes of variation of C, and decide how many modes to use, depending on how we want to trade data reduction against accuracy. (We must also collect data from our M samples that will allow us to iteratively correct our estimates of the object when we are presented with an image containing a new instance of the object. This is explained below.)

An outline of the algorithm that we apply is given here and is based on the work of Cootes[8].

1. Initialize the shape parameter vector.

$$\bar{b} = \bar{0}$$

2. Generate current estimate of object shape.

$$\overline{x_M} = \bar{\mu} + \Phi \bar{b}$$

3. Map current estimate of object shape into an object of this shape with a specific position, orientation, and scale within the given image.

$$\bar{x_I} = m(\overline{x_M})$$

4. Update each point in the current estimate, using information we collected in the sample data A. In general, this update process will change the object position, orientation, scale, and shape.

$$\bar{x'}_I = u(x_I)$$

5. Map our updated object in the image frame back to model space. This gives us the object's shape only, and takes away the information about its position, orientation, and scale within the image.

$$\bar{x'}_M = m^{-1}(\bar{x'}_I)$$

6. Update our shape parameter vector.

$$\bar{b} = \Phi^T (\bar{x'}_M - \bar{\mu})$$

7. Test for convergence, by testing to see if changes in position, orientation, scale, and shape parameters $\bar{b}$ are sufficiently small.

8. If no convergence, return to step 2.

The Step 4 updating of the object position, orientation, scale, and shape in image space is done by moving each landmark point independently along a line segment of length L that is normal to the object boundary at that landmark. At the time landmark data is collected, at each landmark point, grayscale values along such normal segments are also collected to provide a grayscale profile at that point. (More precisely, grayscale differences along the normal are collected, in order to minimize dependencies on absolute grayscale values, and these differences are normalized by the sum of the collected grayscale

differences.) This is done for each landmark point and for each of the M sampled images. If the grayscale profile for a landmark point is modeled as a multivariate Gaussian distribution, the M samples provide data from which the mean and covariance of the profile's distribution may be estimated. For each landmark point, the Gaussian parameters for the profile distribution are calculated. Then, during the ASM search process, grayscale values are also collected from the target image along line segments normal to the current estimate's boundary at each landmark point. These segments are of length L1 > L. For a particular landmark point, ASM searches along this normal for a segment of length L which most closely matches the expected grayscale profile, based on the Gaussian distribution for the profile at that landmark point. The center point of the segment that mostly closely matches becomes the updated landmark point. The metric used to measure closeness of fit between the sampled and expected profiles is the Mahalanobis distance, which may be calculated from the sampled profile, and the parameters (mean and covariance) of the Gaussian distribution that defines the expected profile for that particular landmark point. Figure 3 illustrates the geometry of the grayscale profiles at data collection time; Fig. 4 shows the normals that are used at ASM search time to find the closest matches to the expected profiles. See Cootes[8] for a detailed treatment of the ASM algorithm, as we have used it.

For a comprehensive description of statistical approach to shape analysis, see Dryden and Mardia[9].

5. RESULTS

Expected results in using sample coordinate data relative to the ABCS system to initialize searches for vertebrae. In order to obtain an initial assessment of the feasibility of using sample vertebrae location data (manually acquired) to use for deriving a priori values for iterative algorithms to do vertebral searching, we collected coordinate data on 65 sample images. In each image, the four corner points of the C3 vertebrae were used. The data was collected in a coordinate system fixed to the image frame. We computed the centroids of each of these 65 vertebrae in these image frame coordinates. We then manually positioned ABCS U/V axes within each image and manually collected coordinate data (in the image frame) which defined the U axis and the coordinates of the U/V origin. This ABCS data was used to calculate the U axis inclination angle relative to the image frame X-axis (which, in our image frame system, runs horizontally across the image.) Finally, for each image, the U axis inclination angle and the U/V coordinates in the image frame were used to convert the (x,y) coordinates for the C3 centroid to (u,v) coordinates, measured relative to the ABCS coordinate system for that image. Results are shown in the form of a scatter plot of the centroid points in Fig. 5, where the centroid coordinates are plotted in both the image frame and the ABCS systems, and show, as expected, a clear reduction in variance in the ABCS system as compared to the image frame. Standard deviations for the coordinates are, in pixels, (139,176) in the image frame and (37,49) in ABCS. This suggests that, if the ABCS axes can be located with small errors, the mean coordinate values of sampled vertebrae may be useful candidates for initializing vertebral search algorithms.

Initial ASM results for vertebral location and segmenting. For Test 1 of the ASM method, we collected landmark data on 20 cervical spine images, using vertebrae C3-C5, and the lower half of C2, for our testing. The images were chosen randomly from a set of x-rays taken on subjects 60 years of age or older. For each vertebra, 24 points were collected (11 for C2); six of these correspond to the standard 6-point morphometry set (corners and superior/inferior midpoints on each vertebra). The other points consist of anterior and posterior midpoints and, around each of the corners, four additional points were collected, two on each side of a corner point. These additional corner-area points were collected with the goal of achieving better definition of vertebral curvature in the regions of high curvature change. The study was conducted purely as a proof of engineering concept; medical expertise was not incorporated at this stage of our work. The landmark points should be recognized only as layman's interpretation of vertebral boundaries. Figure 6 shows the statistical model created from the landmarks on these 20 images: for each landmark point collected, the model shown in Fig. 6 gives the mean value for the landmark coordinates. The covariance matrix for this set of landmark data was used to derive the modes of variation shown in Fig. 7. (Only 5 of the modes are shown for purposes of illustration.) The modes show the different ways in which the shape geometry is allowed to change to match shapes found in target images. An illustration of the algorithm in use is shown in Fig. 8 and Fig. 9. In Fig. 8, a test image (on which the model was trained) with known vertebral boundaries is shown, along with the Test 1 model, which has been initialized to an a priori position, orientation, and scale which corresponds to the mean values of those quantities in the sample 20 landmarked images. Figure 9 shows the result in zoomed-in form after convergence. The solid lines show the converged model, and the dots show the target boundary. The convergence shown may be considered quite good. However, this should be expected, since the test image is one the model was trained on; that is, the values of the landmarks for this image contributed to the statistical model used to search the image. A minimum requirement for a reliable model is that it perform consistently with respect to the data on which it was trained. That is, if we use the model to search for one of the objects used to build the model, the ASM algorithm should yield an accurate location for that object. Figure 10 shows the results of performing such a consistency check for Test 1. Each of the images that

contributed to this model was input to the ASM algorithm, and a vertebra search was done for each image. The solid line represents the results of these searches when a priori position, orientation, and scale for the model was set to the mean values of these parameters from the 20 sample images. The vertical axis shows the average point-to-point error in pixels between the converged landmark points and the original landmark points collected on the object's boundary. In only four cases can the errors be reasonably called small, and in 8 of the cases the errors may be considered huge (in the neighborhood of 100 pixels or more). The dashed lines show the results when the model was applied to the same images, but with the a priori position manually set near (in the neighborhood of 50 pixels) the known truth values. The effect of good a priori values for the algorithm is clearly illustrated, although many of the errors remain unacceptably large. This is not unexpected in this trial case, due to the small number of image samples, and the manual technique used to collect the landmark data, which results in many inconsistencies in landmark placement among the objects.

Pitfalls in landmarking shapes. Creating consistent landmarks for shapes within biomedical images is a serious practical problem. Without the assistance of a graphically-based software package, it is prohibitively labor intensive. Even with the assistance of such a package which provides basic capability to mark and record coordinate data from images, along with the associated data required by the ASM (such as graylevel values sampled along segments normal to the curve defined by the landmark points), it may be difficult to collect landmarks which are based on consistent features in the images. A researcher may begin with the intent of landmarking a particular feature (a particular protuberance of the skull, for example), only to discover after examining a number of images that the feature of interest is not present in every image. He then has the choice of omitting such images from the model, which decreases the chances that the model will have any application to those images, or of collecting landmark data based on subjective criteria of what appears to be the best approximation to the desired landmarks. (In the later case, the model may apply well to all of the images in the sense that it do a good job in fitting curves to them; but the location of the individual points along the curve as landmarks of particular features becomes less reliable.) Or he may wish to mark key features on a boundary in the image and then mark equi-spaced points between these features, which lie on the same boundary. Marking the points by using only visual judgement for equi-spacing introduces inconsistencies in the location of the individual landmarks.

We conducted Test 2 in which we concentrated only on vertebrae C2 and C3. For this test, we collected landmark data from 65 cervical spine images, using as part of the data coordinate values collected under expert medical supervision. On vertebra C2 the two bottom corners and the bottom midpoint were expertly collected, and on C3, the four vertebral corners, inferior and superior midpoints, and anterior midpoint were collected this way. To this base set of points, we applied linear interpolation to get a total of 64 boundary points for C3 and 16 for C2. The model based on the mean values of these points is shown in Fig. 11. Figure 12 shows 5 of the modes of variation that resulted from the principal component analysis of this set of data. Using linear interpolation between each pair of points given by the medical experts results in a polygonal model for C3 (and a piecewise linear model for C2). A comparison of the shapes of C2/C3 in Fig. 6 and Fig. 11, and also in Fig. 7 and Fig. 12, illustrate the linearity of the Test 2 model. Obtaining landmarks by linear interpolation was done for simplicity of implementation for the purpose of getting first order results. We then performed a consistency check by searching 20 test images for C2/C3, where the test images used were part of the set used to build the model. Results are shown in Fig. 13, both for the case of using sample mean values for a priori position, scale, and orientation, and for the case where good a priori values were used for position. Again, good a priori position appears to be significant in reducing estimate error, and again, large errors are present in many of the test images, which is not surprising, given the fact that we modeled vertebral shape by simple linear segments. We are continuing to explore this problem with more careful landmarking and more extensive data collection.

The ASM work we have done was carried out on a 300 MHz Windows NT machine using the Active Shape Model Toolkit[10] for MATLAB 5.3.

6. CONCLUSION

Computer-assisted indexing of image content for digitized x-ray images of the spine remains a challenging area of image processing research. The field of statistical shape analysis provides a potential area of attaining at least partial assistance from computerized methods for the collection of important geometric image information for biomedical research purposes.

ACKNOWLEDGMENTS

REFERENCES

1. Taylor CJ, Cootes TF, Lanitis A, Edwards G, Smyth P, Kotcheff ACW, " Model-based interpretation of complex and variable images," *Philosophical Transactions of Royal Society of London. Series B: Biological Sciences*, August 29, 1997, vol. **352**, no. 1358, pp. 1267-74.
2. Smyth PP, Taylor CJ, Adams JE, "Vertebral shape: automatic measurements with active shape models," *Radiology*, May 1999, vol. **211**, no. 2, pp. 571-578.
3. Gardner JC, Heyano SL, Yaffe LG, von Ingersleben G, Chesnut CH III, " A semi-automated computerized system for fracture assessment of spinal x-ray films," *Proc. SPIE Medical Imaging 1996: Image Processing*. vol. **2710**, Newport Beach, CA, February 12-15, 1996, pp. 996-1008.
4. Long LR, Thoma GR, "Segmentation and feature extraction of cervical spine x-ray images," *Proc. SPIE Medical Imaging 1999: Image Processing*. vol. **3661**, San Diego, CA, February 20-26, 1999, pp. 1037-1046.
5. Long LR, Thoma GR, "Indexing of image content in spine x-rays," *Proc. SPIE Storage and Retrieval for Image and Video Databases VIII*, vol. **3972**, San Jose, CA, January 23-28, 2000, pp. 55-63.
6. Taveras JM, Morello F. *Normal Neuroradiology and Atlas of the Skull, Sinuses, and Facial Bones*. Year Book Medical Publishers, Inc. Chicago, 1979.
7. Chow CK, Kaneko T, " Automated Boundary Detection of the Left Ventricle from Cineangiograms," *Computers and Biomedical Research*. vol. **5**, no. 4. August 1972, pp. 388-410.
8. Cootes TF, Taylor CJ, "Statistical models of appearance for computer vision," Wolfson Image Analysis Unit, Imaging Science and Biomedical Engineering, University of Manchester, September 13, 1999. Draft report, available at http://www.wiau.man.ac.uk.
9. Dryden IL, Mardia KV, *Statistical Shape Analysis*. John Wiley & Sons, New York, 1998.
10. Active Shape Model Toolkit, v. 1.0, Visual Automation Ltd., Manchester, UK. Information available at http://www.wiau.man.ac.uk.

INPUT	PROCESSING	OUTPUT
(Image)	(Indexing algorithm)	(Indexing Information)

INITIAL VERTEBRAE LOCATION
-Anatomy Based Coords
-Stats of manually-acquired vertebrae locations

VERTEBRAE BOUNDARY
LOCATION
-Active Shapes Modeling

VERTEBRAE MEASUREMENTS

VERTEBRAE CLASSIFICATION

Geometric Measurements
Anterior/posterior heights
Boundary length
Disc spacing

Biomedical classification
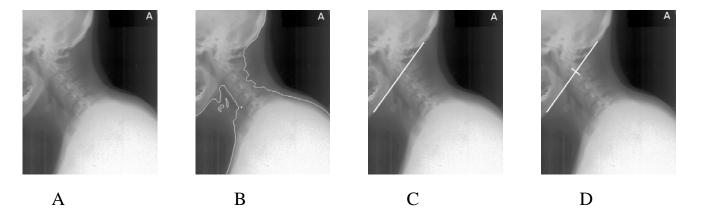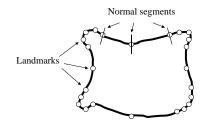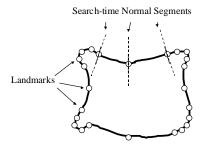Fracture
Anterior osteophytes
Disk space narrowing

Figure 1. Indexing algorithm concepts.

A   B   C   D

Figure 2. (A) Original cervical spine image with borders removed. (B) Foreground/background separation, by grayscale thresholding. The edge of the thresholded image is superimposed on the original image. (C) U axis fixed along the base of the skull. (D) U and V axes fixed. V is shown as the short axis orthogonal to U.

Data collection time: grayscale data is collected along normal segments for each sample object; we get an *expected grayscale profile* for each landmark

Normal segments

Landmarks

One of the sample objects used to build model

Search time: for each landmark point, find the point along its normal segment where the *observed grayscale profile* best matches the *expected grayscale profile* for this landmark

Search-time Normal Segments

Landmarks

Current estimate of object

Figure 3. At data collection time, a profile of grayscale values is collected at each landmark point, along a normal segment of length L. For each landmark point, this collected data is used to estimate the parameters of a multivariate Gaussian distribution for the grayscale profile at that point.

Figure 4. At search time, grayscale profiles of length L are collected at each landmark point, along a normal segment of length $L1 > L$. For each landmark point, the closest match to the expected profile of length L is determined by exhaustively searching the possible subsegments on the L1-length normals. The center position of this closest match determines the updated landmark point. Mahalanobis distance is used to measure closeness between observed and expected grayscale profile.
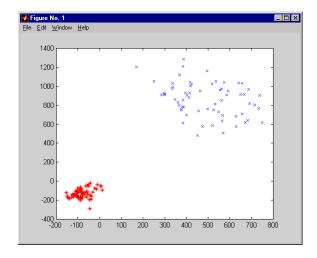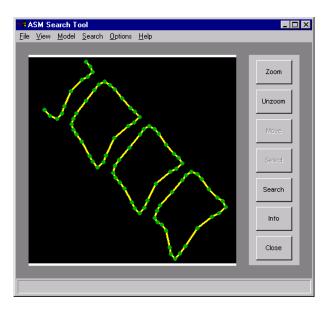

Figure 5. Coordinates of centroids for vertebra C3, for 65 sample images. Upper right shows the coordinates in the image frame system; lower left shows the coordinates in the Anatomy Based Coordinate System (ABCS). Standard deviations for x and y are (139, 176) in image frame and (37,49) in ABCS, in pixels.
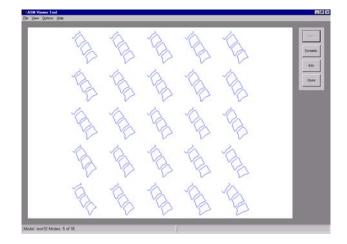
Figure 6. Model for Test 1. This model has 24 landmarks on vertebrae C3-C5 and 11 on C2. (The vertebrae are numbered, from the top, C2-C5.) The model landmarks are the mean values of the landmarks from 20 cervical spine images.

Figure 7. Five of the modes of variation for Test 1. Row 1 is mode 1, row 2 is mode 2, etc. The middle column is the model itself (mean of the 20 samples); the other columns show the variations along the modes at + or − 1 or 2 standard deviations from the mean. For example, entry (row 2, col 4) is +1 s.d. along mode 2.

Figure 8. Test 1 Model applied to a test image, a priori state. The target vertebral boundary is shown on the vertebrae. The starting outline for the model is shown at the mean position, orientation, scale for the sample data.
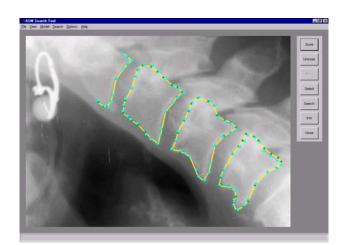

Figure 9. The model in Fig. 8 after convergence. The solid lines represent the converged model. The dots are the target boundary location. In this case, the convergence was accurate.
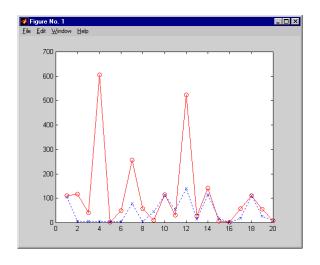
Figure 10. Test 1 model applied to 20 test images. Solid line represents errors from known truth using sample mean values for a priori position, orientation, and scale. Dotted line represents errors when accurate a priori values were specified.
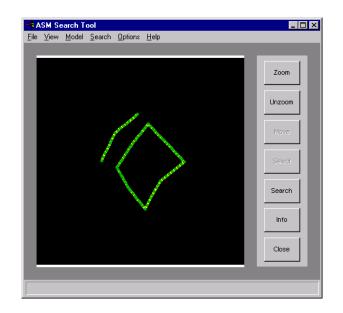

Figure 11. Test 2 model. 65 cervical spine images were used. Only the bottom of C2 and vertebra C3 were included, 64 landmarks on C3 and 16 on C2. A medical expert supplied three C2 landmarks and seven C3 landmarks. The other points were derived by linear interpolation.
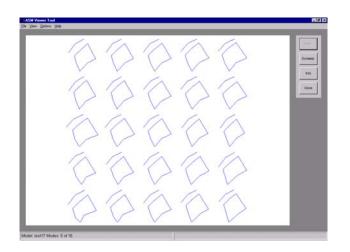

Figure 12. Five of the modes of variation for Test 2. Row 1 is mode 1, row 2 is mode 2, etc. The middle column is the model itself (mean of the 65 samples); the other columns show the variations along the modes at + or − 1 or 2 standard deviations from the mean. For example, entry (row 2, col 4) is +1 s.d. along mode 2.
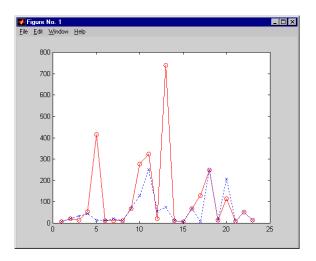

Figure 13. Test 2 Model applied to 20 test images. Solid line represents errors from known truth using sample mean values for a priori position, orientation, and scale. Dotted line represents errors when accurate a priori values were specified.