Identification and classification of spine vertebrae by automated methods

L. Rodney Long[*], George R. Thoma[**]
National Library of Medicine
Bethesda, MD 20894

ABSTRACT

We are currently working toward developing computer-assisted methods for the indexing of a collection of 17,000 digitized x-ray images by biomedical content. These images were collected as part of a nationwide health survey and form a research resource for osteoarthitis and bone morphometry. This task requires the development of algorithms to robustly analyze the x-ray contents for key landmarks, to segment the vertebral bodies, to accurately measure geometric features of the individual vertebrae and inter-vertebral areas, and to classify the spine anatomy into normal or abnormal classes for conditions of interest, including anterior osteophytes and disc space narrowing. Subtasks of this work have been created and divided among collaborators. In this paper, we provide a technical description of the overall task, report on progress made by collaborators, and provide the most recent results of our own research into obtaining first-order location of the spine region of interest by automated methods. We are currently concentrating on images of the cervical spine, but will expand the work to include the lumbar spine as well. Development of successful image processing techniques for computer-assisted indexing of medical image collections is expected to have a significant impact within the medical research and patient care systems.

Keywords**:** x-ray, digital imaging, ASM, active shape modeling, segmentation, classification, NLM, NIH, content-based indexing

1.   OVERVIEW

The goal of our work is the development of computer-assisted methods for the efficient extraction of biomedical content from digitized x-ray images. We have the immediate, specific goal of using this information to index a collection of 17,000 x-ray images of the cervical spine and lumbar spine, and the broader goal of developing methods to index similar collections of x-rays. In view of the increasing sizes of the image repositories being created at medical centers, and the need to utilize the contents of these images with accurate and economical methods, the need for computerized or computer-assisted methods for analyzing, understanding, indexing, and retrieving such images by their content has become a focus of much research effort. A comprehensive summary of the state of the art for developing such systems for general image repositories has been given recently by Smeulders[1].

In the case of our specific image collection, conditions related to osteoarthritis have been identified by a National Institutes of Health (NIH) workshop of rheumatology experts as observable and of interest. These are, for the cervical spine, anterior osteophytes, disc space narrowing, and subluxation; and, for the lumbar spine, anterior osteophytes and subluxation. One goal we have is to classify our images as normal or abnormal for each of these conditions, using training data provided by medical experts. A second goal is to extract general geometric information on the vertebral and spine structure; this information, such as anterior, posterior, and midpoint vertebral heights, and inter-vertebral spacing distance, is expected to be of use in the classification work, but also to have uses beyond this. For example, anterior/posterior height ratios are studied in bone morphometry as a correlate of vertebral fracture. We envision the eventual home for the data that we derive to be in a multimedia database system such as the National Library of Medicine (NLM) Web-based Medical Information Retrieval System[2] (WebMIRS).

Figure 1A provides our concept of the overall technical task. This conceptual framework has evolved from a good deal of experimental work that has achieved success in subtasks of the work for small test sets of images, and is subject to continuing modification as required by future results. Perhaps the simplest description of the task would be "segment, identify, classify", where the primary segmentation is of the individual vertebrae, "identification" refers to labeling anatomy by its specific name (such as "C1" for the appropriate vertebra), and "classification" refers to labeling anatomy as "normal" or

---

[*] long@nlm.nih.gov; phone 1 301 435-3208; fax 1 301 402-0341; http://archive.nlm.nih.gov; National Library of Medicine, 8600 Rockville Pike, Bethesda, MD 20894;
[**] thoma@nlm.nih.gov, phone 1 301 496-4496; other information as above

"abnormal" for one of the conditions of interest (such as "abnormal for anterior osteophytes"). Figure 1A shows this process in an implementable scheme of five "main-line" processing steps and two "off-line" steps.

The first main-line step is the acquisition of basic landmark or orientation information from the images. What minimal information can we compute from the image that is accurate enough for later processing steps and can be robustly computed across many images? By experimental smearing and subsampling using a variety of Gaussian filter strengths and subsampling levels, we observed that even with very heavy smearing and subsampling, three "blobs", distinguished by their brightness characteristics, remain prominent in the images. Two of these are bright areas that correspond to skull and shoulder regions; the third is dark and corresponds to the back-of-spine background. This observation motivated the development of an algorithm to obtain landmarks for skull, shoulder, and (back-of-spine) background in the full size images by analyzing grayscale blobs in heavily smeared, subsampled versions of these images. One of the images with the basic landmark labeling is shown in Figure 1B, upper left. This labeling work is described in a previous paper[3] and, also, in summary form, in Section 2.2 of this paper. The second step shown in Figure 1A is the use of this basic image landmark information to derive a first approximation to the spine region of interest (ROI), by optimizing the location of "reference curves" along the spine in the smeared images, and then using these reference curves to position a spine template in the image. (The spine template is computed in one of the off-line steps as the statistical mean of manually acquired spine ROI boundaries in a number of smeared spine images.) Output from the second step is not only the positioned spine template, illustrated by the second image (clockwise, from upper left) in Figure 1B, but also vertebrae counts, identification, and coarse inter-vertebral boundary marks. Work toward this second step is the main technical result of this paper, and similar work is reported by Zamora[4] and Sari-Sarraf[5]. The third processing step is the iterative improvement of the fitting of the spine ROI template to the image data, using deformable template methods (such as Active Shape Modeling, or ASM). Output is the final, positioned spine ROI. Step four uses the spine ROI just computed, along with the vertebrae counts and inter-vertebral boundary marks, to position a fine resolution template of the vertebrae onto the spine. Image input to this step is the original full resolution image; the template input to this step is pre-computed in the second off-line step, from manually acquired data defining the vertebral boundaries in a sample of the original full resolution images. Output is the set of segmented vertebrae, and is illustrated by the third image (clockwise, from upper left) in Figure 1B. We have previously reported[6] promising work related to this step for small test sets of these images, using manually-acquired vertebral boundary data sampled from 15-20 images and deformed to fit image data by an implementation of the Active Shape Modeling (ASM) algorithm. Work done independently by Sari-Sarraf[5] using similar techniques and tools also appeared to give successful results to a first-order level in a significant number of cases. It should be noted that both in our work and the work of Sari-Sarraf a number of problems were observed and technical issues that require resolution were raised. Among the most outstanding of these are the need for a good method for initializing the ASM step to segment the vertebrae (although Zamora[4], and we, in this paper, have done work toward that goal), the need to investigate the effects of modeling the image grayscale distribution in the neighborhood of vertebrae boundary points with a mixed Gaussian probability distribution function, and the need to understand nonconvergence of the algorithm for certain images even when the template initial conditions are set near known truth. The final indexing and classification step takes the segmented images along with expert training data, and labels the anatomy as classes of normal or abnormal for the conditions of interest. It also generates all of the geometry measures desired from the data. Work toward this step that applies radius of curvature criteria to segmented vertebrae boundaries has been reported by Stanley[7].

As a training set for the classification work, we have collected 550 images of the cervical and lumbar spines (275 cervical, 275 lumbar); all have been classified by a board-certified radiologist as having anterior osteophytes either present or absent. For classification of disc space narrowing and subluxation, we currently have only a few medically validated examples. The ready availability of good training data for anterior osteophytes therefore motivates us to first focus on completing a system that will classify vertebrae for that condition.

2. ANALYSIS

2.1 Hierarchical image analysis and segmentation

Our work in developing automated methods for analyzing the content of these images assumes a hierarchical approach of decomposing the images into structures known at increasingly finer granularities. We can describe this hierarchy as three levels of processing, each with its own set of output data, where each level except the first is dependent on the output from the previous level for its calculation. The outputs at the three levels are, respectively: (1) basic orientation data, based, for robustness, on coarse level features only, providing basic landmarks in the image, with sufficient accuracy to allow initialization of processing for the second level; (2) boundary data for the spine region of interest in the image, defining the spine region at a coarse level, but with sufficient accuracy to allow initialization of the third level; and (3) boundary data for

the individual vertebrae in the spine, calculated with sufficient accuracy to allow vertebral geometry to be used for classification into normal/abnormal categories for specific conditions related to osteoarthritis, e.g. anterior osteophytes, disc space narrowing, and subluxation. Each of these outputs results from a processing stage with its own set of algorithms adapted for the particular processing requirements of that stage. For stage 1, (Figure 1A, step 1) we have used a heuristic analysis of very heavily smeared, highly subsampled images to obtain basic orientation data that we describe in Section 2.2. Stage 2 (Figure 1A, steps 2-3) is still in progress; we are approaching this stage as an optimization problem, with results as given in this paper. Stage 3 (Figure 1A, step 4) is expected to use deformable template processing to locate individual vertebral boundaries at a finely grained level. As noted above, Figure 1B shows images illustrating the expected outputs from step 1, steps 2-3, and step 4, respectively.

The middle image in Figure 1B shows a cervical spine image that has been heavily smeared through Gaussian filtering to obscure details irrelevant to coarse level boundary determination. Superimposed on this image is a template for the spine region of interest, hand-drawn for illustration; this is the type of output we expect to arrive at for steps 2-3. Being able to accurately position such a template on the spine region is expected to allow good initialization of an iterative template-fitting algorithm such as Active Shape Modeling (ASM) for finding individual vertebrae accurately in step 4. We seek to be able to place such a template in these images with a robust algorithm, and first studied image characteristics of the spine region in these images for features that appeared feasible for detection and useful for placing such a template. Figure 2 is one of the x-ray images at full spatial resolution; Figure 3 is a surface plot of this image (reduced in spatial resolution horizontally and vertically by a factor of 64, for the purpose of looking at the large-scale surface characteristics). From Figure 3, some characteristics of the spine region that appear are: (1) the spine is a ridge in the data visually distinguishable from the background, but with less prominence than either the skull or shoulder regions and (2) for most of its extent, the spine has a visually distinguishable slope, on the front and the back, but particularly on the back side. The slope on the front side tends to be less distinguishable for the upper part of the spine, where image data for the jaw is in close proximity, but it is quite prominent at the lower front of the spine, where the image data has an apparent "valley", or at least a point of inflection, that shows up on the image (Figure 2) as a prominent dark area in the throat region. These observations, repeated over some dozens of images, suggested that the following curves may be detectable in the images: (1) a boundary curve for the back edge of the spine, separating spine from the very dark background; (2) a curve following the prominent ridge within the spine itself; (3) a boundary curve following the front edge of the spine, separating the spine from non-spine bright tissue, or from background; and (4) a curve following the prominent valley in front of the spine. Examples of these curves, denoted C_1-C_4, respectively, and hand-drawn for illustration, are shown in Figure 4, superimposed over a smeared version of the Figure 2 image. (This terminology should not be confused with the common anatomical designations for cervical spine vertebrae.)

2.1.1 Detectability of the curves C_1-C_4

Visual inspection of images suggested that the C_1 curve is a very good candidate for a detectable feature, owing to the high contrast, running essentially the entire length of the spine, between the back-of-spine background, and the spine itself. It should be noted, of course, that detecting C_1 in the smeared images yields a curve that would not be expected to mark a significant part of the boundary of any anatomical object in the original unsmeared images, since the original images show the protruding, separated bodies of the spinous processes on the back of the spine, while these bodies form a largely homogeneous, bright mass in the smeared images. At best, C_1 as detected in the smeared images could be expected to be an envelope curve that is tangent to the extremities of the spinous processes.

Similarly, the C_2 curve appeared to be a good candidate for detection, owing to the visually recognizable ridge characteristics in the observed spines, usually extending the length of the spine, but particularly visible in the region of the lower vertebrae. Note that one definition for curvature of the spine might be taken as the curvature defined by a curve that is fit to the midpoints of the top and bottom of each vertebral body. Unfortunately, these midpoints do not have any prominent associated visual characteristics, either in grayscale or in shape, and are very poor candidates for detection until a later stage of processing when the spine anatomy is already known at a finer grained level. However, the grayscale ridge points, i.e., the points on the C_2 curve, appear to lie on the vertebral faces, or near the right edges of the vertebral faces, and hence C_2 might be conjectured to have approximately the same curvature as a curve fit to the vertebral upper and lower midpoints. Thus the curvature of C_2 might be conjectured to give a reasonable approximation to the spine curvature.

Curves C_3 and C_4 are more problematic, since neither the front edges of the vertebrae nor the dark valley in front of the vertebrae are prominent features along the entire spine length. In an original resolution image, it is frequently observed that the front vertebral edge adjoins a region of tissue that appears to be a relatively bright, narrow strip that merges with this edge. Ideally, we would be able to detect a curve C_3 that would mark the boundary between the vertebrae and this tissue; and, in fact, this is a goal of the final stage of the segmentation of these vertebrae. However, in the smeared images, the

vertebrae front edges and the background tissue are no longer distinct. C_3 for these images may be expected to be a curve that follows the general shape of the front of the spine, with best agreement with that shape for the lower spine vertebrae area. Likewise, the valley used to define C_4 in the smeared images yields only an approximation of the valley curve in the original image, and, again, the greatest agreement with the general shape of this curve in the original images is expected in the lower vertebrae regions.

The use of this curve data, if successfully and robustly derived, is to allow placement of a pre-computed template of the spine region. This template is computed as a statistical mean of manually sampled spine region boundaries. It is not known a priori which of the individual C_i or which combinations of the C_i will be most useful in placing the template. In principle, the ability to reliably compute only two points on the template in image coordinates (one point to fix the template on the image up to rotation—the other, to fix the rotation angle) is adequate. But we expect the practical solution of template placement to require the computation of an overdetermined solution using noisy C_i curve data and statistical methods. Figure 5 illustrates these idealized C_i curves that are discussed in more detail in Section 2.3.

2.2 Previous results in obtaining basic landmarks of image regions
For a test set of these x-ray images we have previously produced analysis results[3] that appeared, by visual inspection, to provide elementary region identification to a coarse level, for the skull, shoulder, and background regions. These results held for 46 of the 48 test images. The method used was based on analyzing the images for prominent grayscale regions after the images were reduced to extremely low resolutions, with the expectation that the very bright skull and shoulder, and the very dark background regions, would survive the resolution reduction. The images were reduced in spatial resolution in a two-step process of Gaussian filtering, with the Gaussian filter coefficients set for very heavy smearing, followed by subsampling to bring the original image size down by a factor of 2^8 in both the horizontal and vertical dimensions. The two failure cases corresponded to images which both had very strong light leakage along the bottom border, creating a very wide, bright strip across the bottom of the image that contributed a third (in additional to the skull and shoulder) very bright region to the image and confounded the discrimination of skull and shoulder.

As noted above, Figure 2 shows an original cervical spine 1462x1755 image in our test set. Figure 6 shows the resulting 6x7 image after Gaussian filtering and spatial subsampling in both the horizontal and vertical directions. In this heavily smeared and reduced image the bright regions corresponding to the skull and shoulder are detectable by simple grayscale thresholding and connectivity algorithms, using empirically derived values for the parameters for the thresholding and connectivity. Discovering such parameter values that hold for more than a few images is typically a source of major problems in the full resolution images, but in these small subimages where all but the coarsest features have been obliterated, finding the parameter values that worked for the entire image set required only minor experimentation. The upper left image in Figure 1B shows the image of Figure 2 and 6 with regions labeled as skull, shoulder, and background ("SK","SH","BG", respectively). As a region (e.g., skull) was identified in the Figure 6 subimage, the center of gravity (c.o.g.) of that region was computed; this c.o.g. was then mapped back to a point in the full resolution image, and the appropriate label (e.g., "skull") was placed at that point in the full resolution image, as shown in the upper left image in Figure 1B.

2.3 The general optimization problem
In this paper we build upon these previous results to derive estimates for a cervical spine axis line approximating the position and orientation of the cervical spine in the images. Figure 7 shows the results of using the skull and shoulder labels if the upper left image in Figure 1B to determine an approximate axis line for the cervical spine. In this illustration, the line computed from these two labels appears quite consistent with the true position and orientation of the spine; however, in the general case, this line was only accurate enough to be used to initialize further processing to compute additional curves in the spine region, as described here.

We formulated the general problem of this paper, the computation of spine region data for placing a spine template, as an optimization problem in the location of four curves, C_i, i=1,4, defined as follows: C_1: boundary between back of spine and background; C_2: ridge of brightest grayscale points within spine, running from top to bottom of spine; C_3: boundary formed by front vertebral edges and tissue or background in front of vertebrae; C_4: curve formed by darkest grayscale points in "valley" or point of inflection area in front of spine. Terminology: denote our image by I=I(x,y); then, for each i, let C_i be parameterized by t, and let D_i denote the feasible solution region for $C_i=(x_i(t),y_i(t))$; i.e. C_i must satisfy $C_i \subset D_i \subset I$. For each i, D_i is the region in the image plane where, in our formulation, the solution curve C_i must lie and, in the implementation, we constrain the search for C_i to lie within this region D_i. Figure 5 illustrates conceptually the idealized curves C_1-C_4 and their idealized feasible solution domains D_1-D_4. A very general formulation of an objective function J(.) for our problem would have form $J(C_1,C_2,C_3,C_4)$, where J(.) would include coupling between the C_i, and also include constraints to insure that the

solution curves C_i have reasonable geometries for curves representing spine boundaries. The objective would be to minimize $J(.)$ over all curves C_i satisfying $C_i \subset D_i$.

2.4 The limited optimization problem

At the current time, we have not devised a specific formulation for $J(.)$; rather, we posed the more limited problem of finding the optimal C_i individually for each i, by formulating separate objective functions, each with simple form and each decoupled from the others. That is, for each i, we defined an objective function $J_i = J_i(C_i) = J_i((x_i(t), y_i(t)))$. Note that for a given curve C_i in the (x,y) plane, we may denote the corresponding image grayscale profile by $I(C_i) = I(C_i(x_i(t), y_i(t)))$, and the corresponding image grayscale gradient, evaluated along the curve C_i, by $\nabla I(C_i) = \nabla I(C_i(x_i(t), y_i(t)))$. Then, for each i, we seek to minimize the corresponding objective function, i.e., we seek a solution curve $\hat{C}_i$ such that $J_i(\hat{C}_i) \le J_i(C_i)$ for all $C_i \subset D_I$ that satisfy reasonable constraints on the C_i properties.

For the individual i, we defined these specific objective functions:

$$J_1(C_1) = -\int |\nabla I(C_1)| \qquad + \qquad \lambda \int ds \qquad (1)$$

$$J_2(C_2) = -\int I(C_2) \qquad + \qquad \lambda \int ds \qquad (2)$$

$$J_3(C_3) = -\int |\nabla I(C_3)| \qquad + \qquad \lambda \int ds \qquad (3)$$

$$J_4(C_4) = \int |\nabla I(C_4)| + I(C_4) \qquad + \qquad \lambda \int ds \qquad (4)$$

where ds is arc length along C_i, and for each i, C_i is constrained to lie in D_i. It should be noted that our description of the constraints on the C_i omits features that should be incorporated into a more complete model: for example, the C_i should be approximately the same lengths, hence, the begin and end points for each C_i should be constrained to lie within pre-specified regions.

The choice of integrands in equations (1)-(4) was made with both theoretical considerations and experimental observations. Figure 8 illustrates the expected grayscale characteristics along the solution curves C_i. From equations (1)-(4) it can be seen that, for the cases where we are seeking edges (i=1,3), we maximize the absolute gradients; for the cases where we are seeking a ridge (i=2) we maximize the grayscale value itself. For case i=4, we found that for some images, there is a true "valley" (local minimum) in the D_4 region; for these cases, minimizing the grayscale itself ($I(C_4)$) worked well; in other cases, however, lines transverse to the spine that pass through D_4 have monotone decreasing grayscale, with only a point of inflection instead of a valley. For these cases, having the absolute gradient also in the integrand was necessary. For the final formulation, we integrated the sum of absolute gradient and grayscale. In each of the cases, we penalize this minimization by the length of the minimizing curve, in order to avoid solutions where the curve length grows arbitrarily large. The multiplier λ allows for different weighting of the two terms, for example, in order to compensate for different scales of measurement being used for curve length and for gradient measurements.

3. IMPLEMENTATION

For the implementation of the optimization problem for finding the C_i, we have the correspondence in terminology and concepts given in Table 1. Each of the continuous domain minimization problems has an associated discrete domain problem with corresponding discrete objective function, discrete "feasible solution region", and domain of discrete-valued solution curves. We give an overall description of our solution implementation for all of the discrete objective functions J_i^*, and illustrate the J_1^* case in detail.

For solving for each of the C_i^*, the general approach is the same: (1) determine the non-rectangular grid D_i^* which is the discrete version of the feasible solution region; and (2) minimize the objective function $J_i^*(.)$ over all paths on this grid D_i^*. The minimizing path is the desired solution C_i^*. For step 1, the method used is heuristic and requires the use of experimentally determined parameters to determine D_i^* based on expected characteristics of the image grayscale. This is illustrated in Figure 9 for D_1^*. For step 2, the method of dynamic programming is used; this enables the global minimization of J_i^* to be achieved over the grid D_i^* within a reasonable amount of computation. The solutions were computed on the heavily smeared images, but at full spatial resolution (1462x1755). For each image, the grayscale values were initially

normalized to a 0-1.0 scale, corresponding to the minimum and maximum, respectively of the original grayscale values. The images were smeared with a 100x100 Gaussian filter with standard deviation set to 100. The J_i parameter λ of equations (1)-(4) was set to 0.05 to compensate for the different measurement scales between gradient and distance in the images. For each image, a boundary area of 60 pixels was used on all four image sides (left, right, top, bottom); pixels that were within this boundary limit of an edge were not processed. This was done to avoid the frequent problems encountered by the presence of very bright pixels due to light leakage near the edges of the images.

3.1 Step 1: determining the D_i*

The determination of each of the D_i* uses a common geometrical framework. First, we use the results of the work described in Section 2.2 to define a line segment S taken to be the tentative "spine axis". At this stage of the processing, this "spine axis" is only expected to lie near enough to the actual spine, with a similar enough orientation, so that lines transverse to S will cut the actual spine at an approximately perpendicular angle. We obtain S by connecting the points defined by the "skull" and "shoulder" labels of the Section 2.2 outputs. For subsequent processing, it is necessary to know the approximate points of intersection of this line S with the skull and the shoulder on the original unsmeared image. These intersection points S_{sk}, S_{sh} are found by a curve analysis of the grayscale and grayscale derivative values along S. (Specifically we search S in both directions until we find points satisfying experimentally determined criteria for slope and grayscale value.) Figure 7 shows an example of the line segment with the skull and shoulder intersections S_{sk}, S_{sh} on S marked with boxes.

The transverse lines M_k (again, see Figure 9) were then defined by requiring: M_k is normal to S for all k; M_1 intersects S at $S_{sk;}$ M_{kmax} intersects S at $S_{sh;}$ and the M_k are uniformly spaced along S. After some experimentation, k_{max} was taken to be 20, but no attempt has been made to seek an optimal value. The M_k were taken to span the entire width of the image, except for the image border areas noted above.

Figure 9 illustrates how the grid for D_1* was determined. In Figure 9, the lines M_k are seen to be normal to the "spine axis" S that connects the skull and shoulder landmarks. To determine any of the D_i*, we proceed through the lines M_k one by one, determining which points belong to D_i* on the current line before continuing to the next line. When operating on a particular line, the process is the same for each of the D_i*: (1) determine bounds on the interval on the current line expected to contains points of D_i*, and then (2) sample this interval for the points to actually collect for D_i*. For D_1* processing along a particular line we proceed as follows: first we search the grayscale profile for this line, proceeding from the back-of-skull region toward the front-of-skull region, and, in informal terms, "look for the first occurrence of a large increase in grayscale that occurs over a large spatial span"; i.e., we want to find the approximate point where the grayscale begins climbing rapidly from the very dark background values to the relatively bright spine values as we move along the line from the back-of-skull region toward the front-of-skull region. This point is taken to be a rough indicator or landmark of the D_1* region; we then sample an interval on the current line, centered around this landmark point, to get the points from this line for D_1*. The process is illustrated in detail in Figure 9. Similar methods were used to determine the grids D_2*-D_4*. Tables 2-5 give a summary of the heuristics used for each of the D_i*. The description of the heuristics given in Tables 2-5 list specific parameters used in the formulation of the heuristic, and give the value used for these parameters.

3.2 Step 2: solving for the J_i*

The method for optimizing the J_i* is illustrated in Figure 10, for J_1*. The objective function for J_1*, shown in Figure 10, is minimized over all possible paths on the grid D_1*. The top illustration in Figure 10 suggests the combinatorial expense of directly computing J_1* for all paths. By casting this as a dynamic programming problem, however, the computations are made feasible within practical computation time. It should be noted that some of the grids D_i* have a dependence on the optimal solution points for other Dj*, j≠i. For example, on any line, the points in D_2* must lie "to the left" (toward the front of the spine) of the optimal point for D_1* on that line. This is a way of specifying the observable fact that the spine ridge point on a line must lie "to the left" of the back-of-spine edge point and means, in practical terms, that we must solve for the back-of-spine curve C_1* before we solve for the spine ridge curve C_2*, for example. The bottom illustration in Figure 10 shows a particular solution D_1* that has resulted from applying dynamic programming to find the path that minimizes J_1* among all possible paths on the D_1* grid .

4.   RESULTS

We computed the C_i curves for all 46 of the 48 test images for which the basic landmark data of Section 2.2 was obtainable. Figure 11 shows an example of the results of the computation of the curves for one image, overlaid on the original resolution image (although the curves were computed from the smeared image data). Figure 12 is a second example, computed on a different image. We evaluated the results by displaying and informally inspecting each of the four curves for each of the 46

images, and by using the C_2^* curves to estimate orientations of the spine which we compared to an independent data source. From the visual inspection of the results, we concluded that, when overlaid on the original unsmeared images, (1) the C1* curves tended to fall along the faint boundary separating the non-bony tissue along the back of the neck from the x-ray background; (2) the C_2^* curves strongly tended to follow the spine curvature, and tended to lie along the rightmost edge of the vertebra face; (3) the C_3^* curves tended to lie along the vertebrae front edges, for the lower spine vertebrae, but sometimes fell close to the C_2^* curves for the upper vertebrae; (4) similarly, the C_4^* curves tended to fall in the dark tissue area in the throat area for the lower vertebrae, but also tended to fall close to the C_2^* curves for the upper vertebrae. Overall, the visual inspection supported the hypothesis that the C_2^* curves were the most predictably correlated with the spine anatomy over the length of the cervical spine. To obtain a quantifiable performance measure, we computed linear fits to each of the C_2^* curves. Figure 13 illustrates linear fits computed for all of the Ci* curves of Figure 11. For each of the 46 images in our test set we have available manually-collected (x,y) coordinates for key landmark points on the vertebrae, including the vertebrae "corner" points, and the mid-points of the vertebrae tops and bottoms. This data was collected under supervision of a board-certified radiologist with training in bone anatomy. For each image, we also computed a linear fit to the vertebrae top and bottom midpoints for the image. This gave two straight lines for each image, one computed fitting to our C2* curve, and one fit to the manually-collected coordinates. For each of the two straight line, the slope angle, which may be taken as an estimate of spine orientation, was computed and the absolute difference taken. The results are plotted in Figure 14 and show, that, with one exception, the difference in spine orientations between the two methods as being less than 15 degrees. Areas for improvement in this work include not only possible improvement in the heuristics for computing the D_i^*, but perhaps also improvement in the integrands used in the modeling of the objective functions J_i of equations (1)-(4), in particular for the case J_4, where the image grayscale characteristics are more variable and less understood than any the other cases.

Our next steps will be to evaluate the use of the obtained C_i data for spine template placement, vertebrae counting, and obtaining coarse intervertebral boundary marks. All algorithms were implemented in MATLAB 5.3 and executed on a 400 MHz PC.

5.   CONCLUSION

The problem of analyzing a large collection of digitized x-rays efficiently using computerized methods is a longstanding problem that has only obtained piecemeal success. Through the concerted efforts of several research groups, a general, systematic approach is being formulated. Progress has been made in the areas of automating the acquisition of basic landmark information, obtaining approximate spine location and orientation, determining approximate vertebral boundaries using deformable template methods, and classifying vertebral shapes for anterior osteophyte presence. In this paper we have reported specific results for obtaining approximate spine orientation in an automated fashion.

ACKNOWLEDGMENTS

REFERENCES

1. Smeulders AWM, Worring M, Santini S, Gupta A, Jain R. Content-based image retrieval at the end of the early years. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 22, No. 12, December 2000, pp. 1349-1380.
2. Long LR, Pillemer SR, Lawrence RC, Goh G-H, Neve L, Thoma GR. WebMIRS: Web-based Medical Information Retrieval System. *Proceedings of SPIE Storage and Retrieval for Image and Video Databases VI*, SPIE vol. 3312, San Jose, CA, January 24-30, 1998, pp. 392-403.
3. Long LR, Thoma GR. Feature indexing in a database of digitized x-rays. Proceedings of SPIE Electronic Imaging 2001: Storage and Retrieval for Media Databases. Vol. 4315, San Jose, CA, January 21-26, 2001, pp. 393-403.
4. Zamora G, Sari-Sarraf H, Mitra S. Estimation of orientation of cervical vertebrae for segmentation with active shape models. Proceedings of SPIE Medical Imaging 2001: Image Processing, San Diego, CA, February 17-23, 2001 (forthcoming).

5. Sari-Sarraf H, Mitra S, Zamora G, Tezmol A. Customized active shape models for segmentation of cervical and lumbar spine vertebrae, Texas Tech University College of Electrical Engineering technical report, available at http://www.cvial.ttu.edu/~sarrafT

6. Long LR, Thoma GR. Segmentation and image navigation in digitized spine x-rays. Proceedings of SPIE Medical Imaging 2000: Image Processing. Vol. 3979, San Diego, CA, February 12-18, 2000, pp. 169-179.

7. Stanley RJ, Long R. A radius of curvature approach to cervical spine vertebra image analysis. Proceedings of the 38th Annual Rocky Mountain Bioengineering Symposium, Copper Mountain, Colorado, April 20-22, 2001 (forthcoming).

8. Chwialkowski MP, Shile PE, Pfeifer, Parkey RW, Peshock RM. Automated localization and identification of lower spinal anatomy in magnetic resonance images. Computers and Biomedical Research, vol. 24, no. 2, April 1991, pp. 99-117.
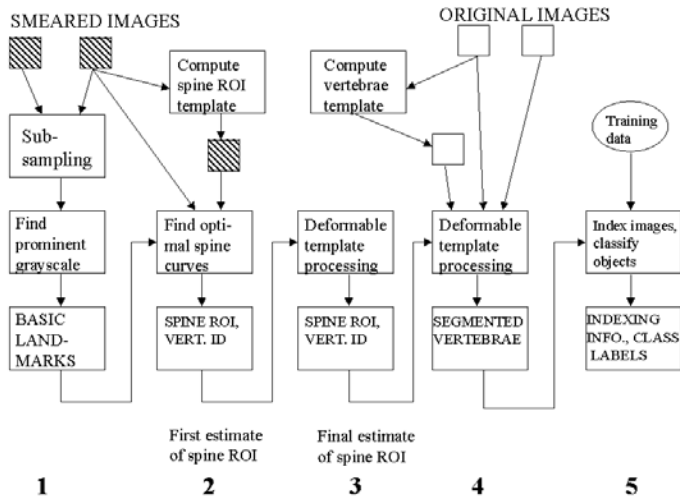
Figure 1A. Conceptual overview of the technical steps in the indexing and classification work. The work is broken into five main processing steps as numbered above, plus two "off-line" steps required to produce templates for the spine ROI and for the vertebrae.
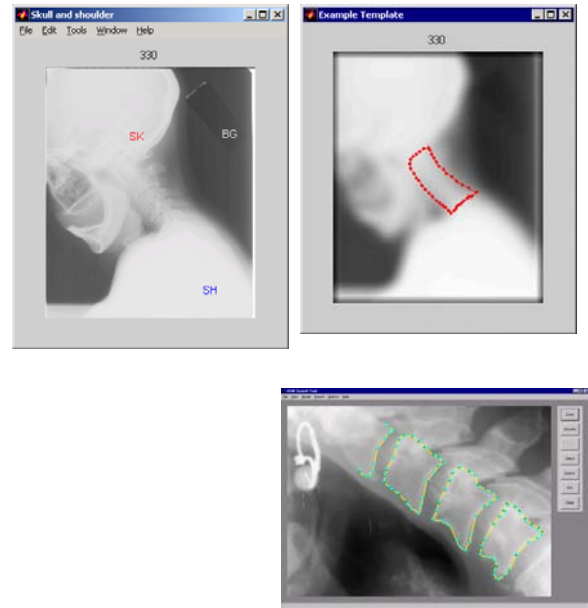
Figure 1B. The images show, clockwise from upper right,, an example of the basic landmark labels (for skull, shoulder, and back-of-spine background) output from step 1; an example of the spine ROI boundary (hand-drawn for illustration) output by steps 2 and 3; and an example of the vertebrae segmentation output by step 4. In this illustration, the step 4 output was produced by using an implementation of the Active Shape Modeling (ASM) algorithm.
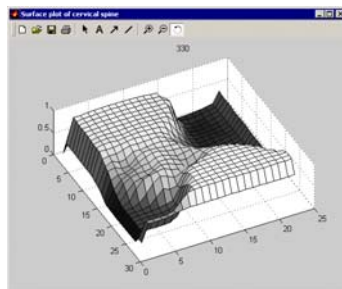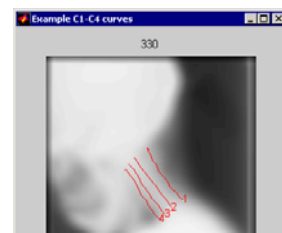

Figure 3. 2.

Figure 4. The four spine curves: C1: back of spine edge, C2: spine grayscale ridge, C3: front of spine edge, C4: grayscale valley
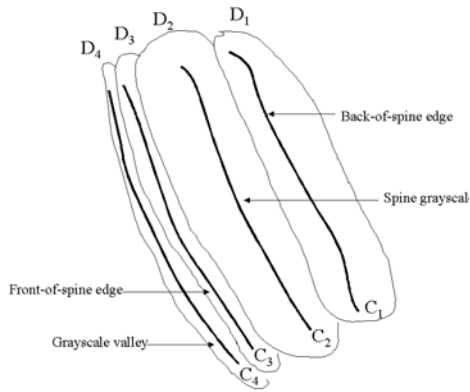
Figure 5. The idealized curves CB_{1B}-CB_{4B} and their (idealized) feasible search domains DB_{1B}-DB_{4B}. In the general optimization problem the solution curve Cj is coupled with the other solution curves CB_{iB},i=1,4, i≠j.
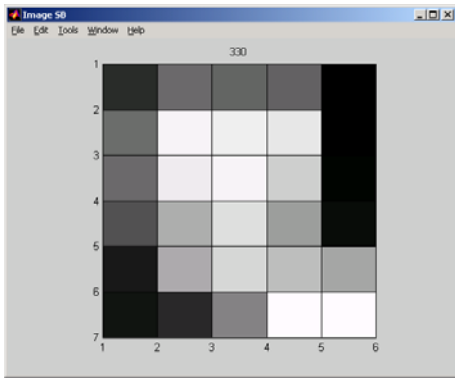

Figure 6. The image from Figure 2, after heavy Gaussian filtering and subsampling. Skull and spine bright areas are visible, and dark background.
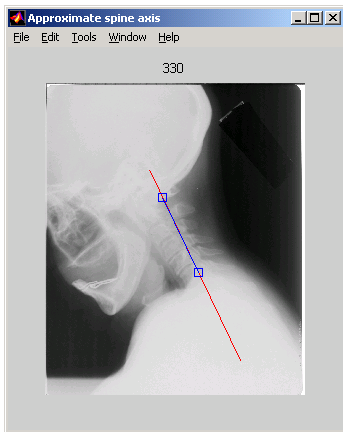

Figure 7. The line S is defined by the SK (skull) and SH (shoulder) points of Figure 6. The boxes mark the intersection points SB_{skB}, SB_{shB} of S with the skull and shoulder boundaries.
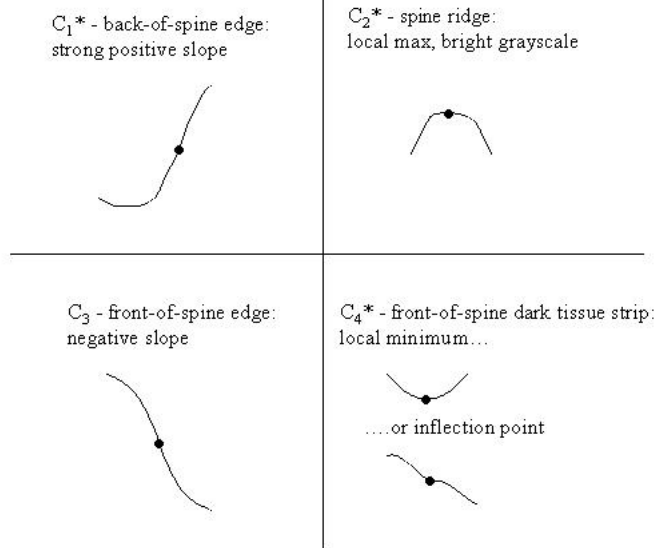
Expected features for points on C_i^*

C_1^* - back-of-spine edge: strong positive slope

C_2^* - spine ridge: local max, bright grayscale

C_3 - front-of-spine edge: negative slope

C_4^* - front-of-spine dark tissue strip: local minimum…

….or inflection point

Figure 8. Expected features for points on the curves CB_{iB}^*.

Grayscale profile $I(\mathbf{q})=I(x(t),y(t))$

Skull landmark

The lines M_k

Shoulder landmark

Step 1: On each line, the first point of large grayscale increase is determined by finding the first $\mathbf{p}$ such that $I(\mathbf{p}-\Delta_2*\mathbf{uhat})-I(\mathbf{p})>=f_1*cd_{max}$

Step 2: Then we sample discrete points on the line between $\mathbf{b_1}$ and $\mathbf{b_2}$ to compute the grid points for a line. The non-rectangular grid D_1^* is the set of all such grid points for all of the lines M_k.

Symbols
Δ_1 = half-length of search interval
Δ_2 = interval used to detect large grayscale change
cd_{max} = maximum grayscale change over entire profile
f_1 = fraction of cd_{max} used to compute threshold for $\mathbf{p}$
$\mathbf{uhat}$ = unit vector parallel to M_k lines, pointing toward back of spine region

Figure 9. Heuristic for determining DB_{1B}^*

Figure 11. Computed curves CB$_{1B}$*-CB$_{4B}$*, overlaid on image.

Figure 12. Second example of computed curves CB$_{1B}$*-CB$_{4B}$*.
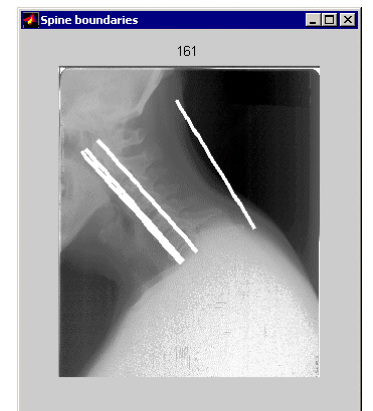
Figure 13. Linear fits to the data of Figure 11.
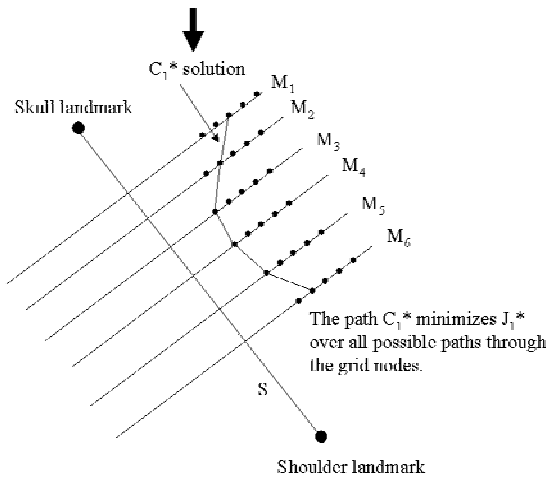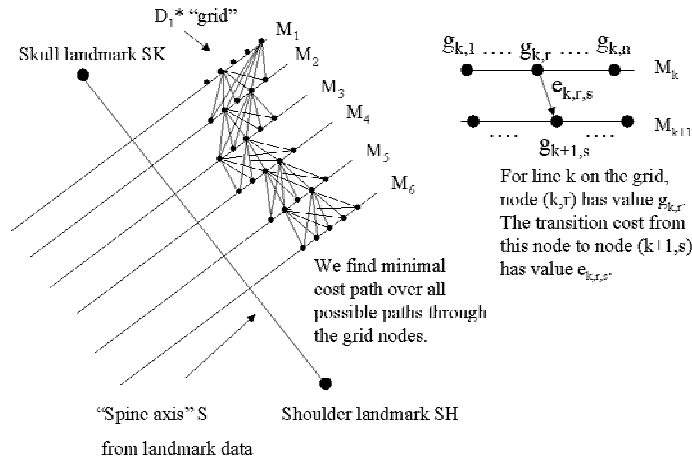

Figure 10. Finding the optimal solutions CB$_{iB}$*. Illustrated for the JB$_{1B}$* case. We implement DB$_{1B}$ as the non-rectangular grid DB$_{1B}$* of nodes placed on the lines MB$_{kB}$ that are normal to the approximate "spine axis". Each node has an assigned value gB$_{k,rB}$ (gradient magnitude) and transition cost eB$_{k,r,sB}$ (Euclidean distance) for moving from node (k,r) to node (k+1,s). To every path C consisting of an ordered set of nodes (g B$_{1,j1B}$, g B$_{2,j2B}$,…,g B$_{kmaxB}$jB$_{kmaxB}$), one per line, beginning at MB$_{1B}$ and ending and MB$_{kmaxB}$, we assign an associated cost:

$$ J_1^*(C) = -\sum_{i=1}^{k\,\max} g_{i,j_i} + \lambda \sum_{i=1}^{k\,\max\,-1} e_{i,j_i,j_{i+1}} $$

We then minimize JB$_{1B}$*(C) over all possible such paths C, as suggested by the figure above. The dynamic programming technique reduces the combinatoric possiblities to a tractable size. We take the minimizing path

CB$_{1B}$* = arg min JB$_{1B}$*(C)

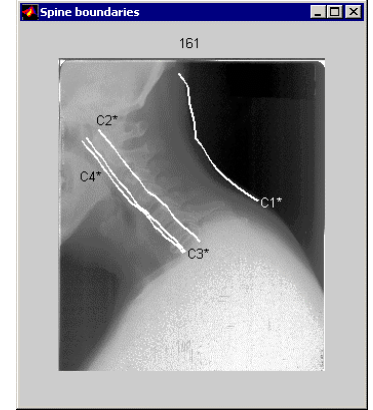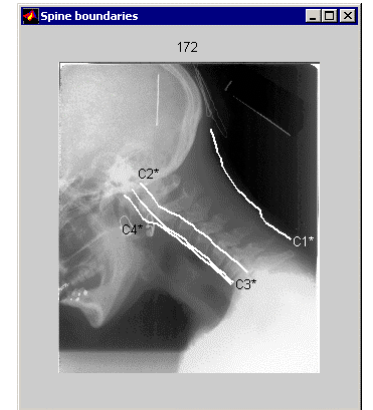to be the discrete, computable solution corresponding to the ideal spine region curve CB$_{1B}$.
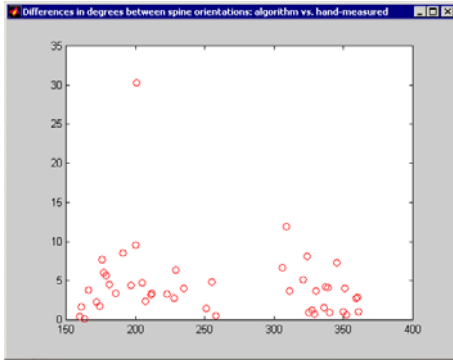
Figure 14. Differences (in degrees) between estimates from spine orientations derived by using (1) linear fits to the C2* curves vs. (2) linear fits to manually-collected coordinate data for the spines.

	Continuous Domain	Discrete Domain
Objective Functions	$JB_{iB}(.)$ (sums of integrals)	$JB_{iB}*(.)$ (discrete sums)
Feasible regions	DB_{iB} (regions in the plane)	$DB_{iB}*$ (non-rectangular grids)
Solution curves	CB_{iB} (smooth curves)	$CB_{iB}*$ (piecewise linear curves)
Solution domain	All curves in DB_{iB} satisfying reasonableness criteria for spine curves	All paths on the grid $DB_{iB}*$
Optimization	--	Dynamic programming

$DB_{1B}*$		
Curve desired	Back of spine edge	
Characteristic feature of curve	Strong positive directional derivative	
How grid bounds determined on a line MB_{kB}	Interval of length $2\Delta_2$ centered around landmark point **p**	
Search region for **p**	First half of the line MB_{kB}	
Criteria for **p**	First **p** so that $T(\mathbf{p}) \equiv I(\mathbf{p} - \Delta_1\hat{\mathbf{u}}) - I(\mathbf{p}) \geq f_1 d_{max}$	
Δ_1	Distance over which to look for large grayscale increase	150
Δ_2	Half length of interval size for $DB_{1B}*$ on line MB_{kB}	100
f_1	Fraction of maximum grayscale profile change for threshold computation	.25
d_{max}	Maximum value of T(**p**) over the entire profile of the line MB_{kB}	
$\hat{\mathbf{u}}$	Unit vector parallel to line Mk, oriented in direction toward back-of-skull region	
p	Point on line MB_{kB}. $\mathbf{p} \equiv (x(t), y(t))$, where t parameterizes the line MB_{kB}	
I(**p**)	Image grayscale function, evaluated at **p**	

Table 2. Summary of heuristics for determining $DB_{1B}*$, with parameter values used. Compare Figure 9.

$DB_{2B}*$		
Curve desired	Spine ridge	
Characteristic feature of curve	Local maximum, strong grayscale	
How grid bounds determined on a line MB_{kB}	Interval of length $2\Delta_2$ centered around landmark point **p**	
Search region for **p**	Interval bounded by optimal point for $DB_{1B}*$ on line MB_{kB}, and last point on MB_{kB} line (**)	
Criteria for **p**	First **p** so that $I(\mathbf{p})\, is\; local \max\; and\; I(\mathbf{p}) \geq i2_{tol}$	
$i2_{tol}$	Grayscale threshold that I(p) must satisfy	.5
Δ_2	Half length of interval size for $DB_{1B}*$ on line MB_{kB}	100

Table 3. Summary of heuristics for determining $DB_{2B}*$, with parameter values used. (**) Note that computation of $DB_{2B}*$ requires that optimal points on $DB_{1B}*$ be computed first.

DB$_{3B}$*	
Curve desired	Front of spine edge
Characteristic feature of curve	Negative directional derivative
How grid bounds determined on a line	Interval bounded by optimal points for DB$_{2B}$* and DB$_{4B}$* on line MB$_{k\ (**)B}$

Table 4. Summary of heuristics for determining DB$_{3B}$*. (**) Note that computing DB$_{3B}$* requires that optimal points for DB$_{2B}$* and DB$_{4B}$* be computed first.

DB$_{4B}$*		
Curve desired	Front of spine dark tissue strip: grayscale "valley" or inflection point	
Characteristic feature of curve	Grayscale local minimum or inflection point, grayscale must be clearly below spine ridge grayscale	
How grid bounds determined on a line	Interval of length $2\Delta_2$ centered around landmark point **p**	
Search region for **p**	Let $\mathbf{p}_2^* \equiv optimal\ pt\ for\ D_2^*\ on\ M_k$ Then the search interval for p is bounded by the points $\mathbf{p}_2^* - \Delta_3 \hat{\mathbf{u}}$ and $\mathbf{p}_2^* - \Delta_4 \hat{\mathbf{u}}$, i.e. points on line MB$_{kB}$ lying at distances Δ_3, Δ_4 from $\mathbf{p}_2^*$, both in the direction of the front-of-spine regions (opposite to the direction of $\hat{\mathbf{u}}$). (**)	
Criteria for **p**	First **p** so that $I(\mathbf{p})\ is\ local\ \min\ and\ I(\mathbf{p}) \leq (1 - i4_{tol})I(p_2^*)$ If such p not found, p is set to the left bound of the search interval, i.e. $\mathbf{p} \equiv \mathbf{p}_2^* - \Delta_3 \hat{\mathbf{u}}$. (**)	
Δ_2	Half length of interval size for DB$_{4B}$* on line Mk.	
Δ_3	We require that the begin boundary for the search interval for **p** to be separated from point $\mathbf{p}_2^*$ along line MB$_{kB}$ by this distance.	20
Δ_4	We require that the end boundary for the search interval for **p** to be separated from point $\mathbf{p}_2^*$ along line MB$_{kB}$ by this distance.	400
$i4_{tol}$	The grayscale at **p** must lie significantly below the grayscale at the optimal spine ridge point $\mathbf{p}_2^*$, for **p** to be a candidate for the DB$_{4B}$* grid. This parameter sets the threshold for the grayscale value that **p** must not exceed.	.05

Table 5. Summary of heuristics for determining DB$_{4B}$*. (**) Note that computing DB$_{4B}$* requires that the optimal points for DB$_{2B}$* be computed first.