A method of content-based image retrieval for a spinal x-ray image database

Daniel M. Krainak[*a], L. Rodney Long[**b], George R. Thoma[b]

[a]Department of Biomedical Engineering, The Catholic University of America, Washington, DC;
[b]National Library of Medicine, Bethesda, MD

ABSTRACT

The Lister Hill National Center for Biomedical Communications, a research and development division of the National Library of Medicine (NLM). maintains a digital archive of 17,000 cervical and lumbar spine images collected in the second National Health and Nutrition Examination Survey (NHANES II) conducted by the National Center for Health Statistics (NCHS). Classification of the images for the osteoarthritis research community has been a long-standing goal of researchers at the NLM, collaborators at NCHS, and the National Institute of Arthritis and Musculoskeletal and Skin Diseases (NIAMS), and capability to retrieve images based on geometric characteristics of the vertebral bodies is of interest to the vertebral morphometry community. Automated or computer-assisted classification and retrieval methods are highly desirable to offset the high cost of manual classification and manipulation by medical experts. We implemented a prototype system for a database of 118 spine x-rays and health survey text data related to these x-rays. The system supports conventional text retrieval, as well as retrieval based on shape similarity to a user-supplied vertebral image or sketch.

Keywords: Content-based image retrieval, spine, x-ray, vertebra, digital image, NLM, NCHS, NIAMS, NHANES

1. INTRODUCTION

Digital vertebral morphometry involves using information obtained from digitized lateral spine radiographs to assess vertebral deformities[1]. Measurements such as anterior height (Ha), posterior height (Hp), the ratio of anterior to posterior height (Ha/Hp), the ratio of posterior and anterior heights of adjacent vertebra, the ratio of mid height to posterior height (Hm/Hp), and others have been used in algorithms that define vertebral deformity[2, 3]. In fact, to help resolve inter-radiologist and intra-radiologist differences, some have proposed using more objective measures such as vertebral dimensions as a diagnostic aid[4]. Additionally, determining the presence of vertebral fractures is difficult, but important to both epidemiology and clinical studies, and the ratio of anterior to posterior height has been used as an indicator of fracture[2]. Several competing algorithms have been developed to assess vertebral deformity, and further studies are necessary to evaluate these and other algorithms that may be developed in the future[1].

The National Library of Medicine (NLM) maintains an archive of approximately 17,000 digitized lumbar and cervical spine x-ray images and accompanying data collected in the second National Health and Nutrition Examination Survey (NHANES II). Researchers at the NLM, collaborators at the National Center for Health Statistics (NCHS), and collaborators at the National Institute of Arthritis and Musculoskeletal and Skin Diseases (NIAMS) wish to classify vertebrae from the collected images as "normal" or "abnormal" for conditions of interest to the osteoarthritis community and identified in two NIAMS workshops. In addition, published research in the vertebral morphometry community supports the interest in and need for good algorithms to derive quantitative vertebral information from digitized x-ray images[5]. A major factor driving the development of computer-assisted image analysis is the prohibitive costs of having the images analyzed by a radiologist, and thus, a computer-assisted method of classification provides a highly desirable alternative.

*74krainak@cua.edu; phone 202-238-1526; http://www.cua.edu/; Biomedical Engineering Department, Pangborn Hall rm 131, Catholic University, Washington, DC 20064; **long@nlm.nih.gov; phone 301 435-3208; fax 301 402-0341; http://archive.nln.nih.gov/;

We demonstrate some of the capabilities of a computer system to determine vertebral morphometry measurements such as anterior height (Ha), posterior height (Hp), the ratio of anterior to posterior heights (Ha:Hp), and others from digitized spine x-ray images.   We also describe the ability of our system to retrieve similar vertebral images based upon vertebral morphometry measurements, vertebral shape, or epidemiological data from the NHANES II surveys. A long-term goal is to enable users to have full access to the database of 17,000 spine images for epidemiology studies, clinical studies, normal value comparisons, and shape matching for normal and deformed vertebrae. We developed a prototype that uses 118 images from 74 participants in the NHANES II survey.

2. METHODS

2.1 Development environment

We implemented the functional code and graphical user interface for our prototype system using MATLAB 6.0 on a Pentium III 600 MHz Windows-based PC. All of the code was written as MATLAB m-files, and the Graphical User Interface was developed using the MATLAB Graphical User Interface Development Environment (GUIDE).

2.2 Determining vertebra characteristics

Our prototype database includes 118 images for which the contents have been indexed under supervision of a board-certified radiologist with expertise in bone images at Georgetown University Medical Center. The indexing consisted of recording the coordinates of up to nine key boundary points on each vertebra, including the six standard morphometric points (corners, plus top and bottom midpoints), the anterior midpoint, and points marking the extremities of anterior osteophytes, if present (Fig. 1). Additionally, the radiologist classified the osteophytes on the lumbar images as "fused" or "not fused". Using these nine points, we calculate and store in the database vertebral characteristics including anterior height (Ha), posterior height (Hp), the ratio of anterior height to posterior height, the distance between adjoining vertebrae, and whether osteophytes exist on a particular vertebra.

Rosol suggests that constraining vertebral height measurement so anterior and posterior height measures are parallel to one another enhances longitudinal reproducibility[5]. Therefore, to maintain parallel measurements, we implemented the following algorithm to determine the vertebral heights. We model the vertebral body as a trapezoid. We fit a first order polynomial to the top three points marked by the radiologists (upper corners and midpoint) and another to the bottom three points on the vertebral boundary, later drawing a perpendicular through the upper line at each corner. The distance measured (in pixels) is the distance from the top line along the perpendicular to the intersection point on the bottom line. A similar trapezoid model is used to measure the distance between vertebrae. In contrast to the single vertebra trapezoid model, the boundaries of the trapezoid used in disc space measurement include a first order polynomial fit to the bottom three points of the upper vertebra, a first order polynomial fit to the top three points of the lower vertebra, and two lines perpendicular to the bottom of the upper vertebra that pass through the corner points as marked by the radiologists.

All the distances calculated by our program measure pixels on the digitized image: no calibration marker was recorded in the images, and no accurate spine-to-film and film-to-focus values are available for the x-ray images, making the determination of absolute (physical) measurements from the images problematic. Differences in magnification caused by varying film-to-focus and spine-to-film distances may result in poor assessment of deformity using current algorithms despite best efforts to correct for optical differences among images[1]. Fortunately, dimensionless measurements such as anterior to posterior height ratio, the percent change in adjacent posterior heights, wedge angle, and predicted posterior height will contribute accurate vertebral morphometry measurements to be used in assessing deformity.

Our image retrieval program permits searches over a range of values for any of the characteristics determined computationally from the nine points labeled by the radiologist: anterior height, posterior height, Ha/Hp, minimum distance between vertebrae, maximum distance between vertebrae, and the distance between any two points located on a single image. Our system also allows for searches on some epidemiological data: person ID number, sex, age, race, height, and weight. Additional relevant information from the NHANES II survey can be added to the program in the future, if desired, for example: locations of back pain, or other back and neck information.

The text-only query screen allows the user to either query the text-based information immediately or select either "Retrieve by Sketch" or "Retrieve by Example" which bring up secondary screens. We included the option to cater the query return either as the information in a table, which may later be saved as a text file or as a collection of images to be viewed by the user. An example query is "Find the images for all females having a cervical spine vertebra with Ha/Hp ratio of at most 0.9" (Fig. 2). Selecting "Retrieve images" returns all the images meeting the query criteria on the left and all the available images on the right to permit side-by-side comparison (Fig. 3). The screen shows some of the text information such as race, sex, or age and information extracted from the vertebra points such as disc spacing and Ha/Hp. On the right hand side, the user may search through the entire database for a specific vertebra of a particular person to compare to the query results.

2.3 Query-by-example and query-by-sketch

One of the novel features of our system is the ability to query vertebra shape by an example or a sketch. To support query-by-example and query-by-sketch in our system, we implement an algorithm capable of quantifying similarity between an input shape and a shape in the database. The method requires a transformation to remove differences in position, scale, and orientation between the two shapes. To achieve this, we use a Procrustes shape-fitting algorithm described by Cootes[6] to bring the shapes to a common frame of reference. An important consideration in any system that retrieves shapes by similarity is the metric used to calculate degree of similarity of two shapes: we calculated shape similarity by summing Euclidean distances between corresponding points.

First, we define the shape to be tested as a subset of the points marked by the radiologist. For query-by-sketch, the user defines the input shape by selecting which of the nine points to include and uses a mouse interface to drag the points into the desired spatial configuration. For query-by-example, the user specifies one of the vertebrae which has the radiologist points, and the system uses the shape defined by these points for its similarity search. For either type of query, images are returned in ranked order of similarity, as determined by the sum of Euclidean distances of corresponding points on the example image and the resulting image, with the total number of retrieved images controllable by the user.

Using the mouse in the query-by-sketch screen, the user can configure the 9 index points arbitrarily to define a vertebral shape to submit to the system for searching. The user has the option of selecting or deselecting any of the nine points as well. The user may only want to sketch a particular corner. For example, a sketch of only points 5, 6, 7, and 9 would let the user search for shapes similar to lower anterior corner sketched. To illustrate the ability of the system to return results similar to a sketch, a user could draw two large anterior osteophytes, both of which are angled toward the anterior face of the vertebra (Fig. 4). The results of the query facilitate a comparison between the user sketch to an actual vertebra. The vertebrae are returned with the most similar vertebra ranked number one and similarity decreasing as ranking number increases. The screen displays the sketch, the vertebra returned, and an image displaying the user sketch (+) over the position of the vertebra points (x) of the vertebra returned (Fig. 5).

On the query-by-example screen, the user specifies the input image, vertebra, and points on the vertebra that are used as a search example (Fig. 6). Note that in our current system, the user may only submit images for which the 9 indexing points are already known; the current system is not capable of deriving these points from an arbitrary input image on the fly. We plan to add features that would not only allow a user to mark the nine points on a vertebra, but possibly automatically or semi-automatically detect the shape of the vertebra. The results for a query searching for lumbar vertebrae having large osteophytes returns a screen similar to the sketch screen, only in the sketch area are the nine points from the example image instead of a sketch (Fig. 7).

The current system supports hybrid text/shape similarity searches by letting the user do a pure text query, then letting the user do a shape similarity query on the results returned from the pure text query.

3. RESULTS

The image retrieval system supports pure text searches, query-by-example similarly searches, or hybrid searches that are combinations of these. The pure text searches may pose queries in terms of the NHANES health survey data or in terms of quantitative data, such as vertebral heights, that we have derived from the images. For pure text searches records are

retrieved by exact match to the query, as in a conventional database system. For example, if a query is made for all females with fused osteophytes, exact matching of database records to the query is done, and one hundred percent of vertebrae with fused osteophytes in females are returned. Note that the text-based information obtained from the images can be used to "search itself". That is, once the quantitative characteristics of the vertebrae are put into the database, we can search this data to find collections of images sharing common vertebral characteristics that may not have been obvious when we were restricted to viewing images to find these vertebral characteristics. For instance, the user can query the database for all vertebrae with a fused lower anterior osteophyte on the lumbar x-ray.

Beyond the typical text-based searches, we provide a mechanism for retrieval by content, based upon features specified on an individual vertebra image. The user may search for particular characteristics of an image without describing them in text, but instead using a sketch, or example image with features similar to the ones desired by the user. The vertebral shape representations in our system have a coarse granularity limited by the fact that we have at most 9 boundary points per vertebral shape. With this limitation, the most prominent shape characteristics are the shapes of the anterior osteophytes, when they are present on the vertebrae. Because of this, the shape-retrieval tests that we have made with the system have relied heavily on distinguishing vertebrae with anterior osteophytes from those without such osteophytes, and on retrieving vertebrae based on osteophyte size or the angle that the osteophyte makes with the vertebral body. Preliminary testing of the system has verified its capability to return visually similar vertebral shapes, given an input shape (either sketched by image example) with large osteophytes, vertebrae with similarly-large osteophytes are returned, ranked by similarity to the input. Given vertebrae with osteophytes that lie at extreme angles to the vertebral body (e.g. an anterior osteophyte on the upper anterior corner that point sharply "back" over the top of the vertebra or sharply "down" toward the anterior of the vertebra), the returned vertebrae are seen to resemble the inputs. (This test was done by sketching the input.) Some tests were run to get a first assessment of using shape queries as a discriminator between vertebrae in the database that have anterior osteophytes and those that do not. For this test, an example image having large anterior osteophytes was presented to the system. A query-by-example for similar images was made, and in the images returned, the radiologist classification of anterior osteophytes present/absent was used to assess how many of the returned images did in fact have anterior osteophytes. In the tests runs, the results showed that in fact all of the returned images had true anterior osteophytes, according to the radiologist classification. (The similarly list for these tests consisted of 25 images. Additional tests were run to see the effects of such queries as the sizes of the osteophytes on the input image decrease. It was observed that doing the query-by-example with small input osteophytes could easily result in some images being returned fairly high in the similarity ranking that in fact did not have osteophytes present, according to the radiologist classification. This happens when other points used in the database vertebrae match closely to the query vertebra except for points 8 and 9, the osteophyte points. Consequently, weighting specific areas of the vertebra, such as osteophytes, at the time of query may provide a better match to the input image. All testing to date has been informal and has been intended to obtain a first-order evaluation of the system performance.

4. DISCUSSION

Our program demonstrates the ability to combine text-based information with image-based information when searching through an image database. Searching with image features such as the distance between vertebrae, and searching by shape can be effectively combined with health survey data such as age and sex. The current database contains 118 images for 74 people in which, for each cervical and lumbar spine vertebra, nine boundary points were labeled by a radiologist. The nine points query-by-example and query-by-sketch methods are highly effective in returning visually similar images, and provide strong motivation for the evaluation of a system for retrieving data consistent with the classifications of biomedical experts, and with incorporation of more detailed shape data into the database.

Extracting image-based information from the NLM's archive of digitized spine x-ray images with minimal user effort is highly desirable due to the prohibitively high cost of having a professional label each individual vertebra. Our preliminary work indicates that without individually measuring every aspect of the vertebra, valuable information may be obtained from a subset of information such as the nine points or other abstraction of the vertebra shape. Methods of automated point placement for the key vertebral morphometry points are under investigation[4]. In the future, we hope to use automated or semi-automated methods to analyze the image content of some of the remaining 16,000 images that were not marked by a radiologist. (Active shape modeling methods are under current investigation for segmenting the

vertebrae.) Using our current nine-point model and other models, we will assess the validity of computer-assisted vertebral morphometry measurements compared to traditional methods of measurement.

In the more immediate future, we plan to continue the evolution of this system to incorporate additional image-derived quantities within our database, such as disc spacing, subluxation measurements, vertebral area, wedge angle, and percent difference in anterior height between adjacent vertebrae (PDAH) [8]. We also plan to extend the shape searching to include shapes represented by multiple vertebrae or the shape at the junction of two vertebrae. Finally, we hope to incorporate into the database detailed vertebral boundary data and to support query by user-defined shape, based on this data. This would allow us to computationally analyze vertebral images with much more detailed information than standard six-point vertebral morphometry, which is shown to be useful, but sometimes results in significant measurement errors[7]. With more detailed shape information available, the representation of the vertebral shapes by compact features, such as coefficients of approximating curves, will contribute to improved performance. Comprehensive characterization of the future system's query results in terms of completeness and precision will be a significant measure of the value of the system. The current system is a small prototype demonstrating basic capabilities of a database system for retrieving biomedical information using text and/or shape similarity searches for vertebrae in digitized x-rays of the cervical and lumbar spines, and is a test bed and development vehicle for the continued evolution of such systems.

REFERENCES

1. K. Weber, et. al., "Measurement imprecision in vertebral morphometry of spinal radiographs obtained in the European Prospective Osteoporosis Study: consequences for the identification of prevalent and incident deformities," *The British Journal of Radiology*," **72**, pp. 957-966, 1999.
2. R. Eastell, S.L. Cedel, H.W. Wahner, et al., "Classification of vertebral fractures," *Journal of Bone and Mineral Research*, **6**, pp. 207-215, 1991.
3. E.V. McCloskey, T.D. Spector, K.S. Eyres, et al., "The assessment of vertebral deformity: a method for use in population studies and clinical trials," *Osteoporosis International*, **3**, pp. 138-147, 1993.
4. P.D. Ross, J.W. Davis, S.E. Robert, and R.D. Wasnich, "Ability of vertebral dimensions from a single radiograph to identify fractures," *Calcified Tissue International*, **51**, pp. 95-99, 1992.
5. M.S. Rosol, et. al., "Vertebral morphometry derived from digital images," *American Journal of Roentgenology*, **167**, pp. 1545-1549, December 1996.
6. T.F. Cootes and C.J. Taylor, *Statistical Models of Appearance for Computer Vision,* pp. 84-86, Wolfson Image Analysis Unit, Imaging Science and Biomedical Engineering, University of Manchester. http://www.wiau.man.ac.uk/~bim/Models/app_model.ps.gz Feb 28, 2001.
7. J.C. Gardner, L.G. Yaffe, Johansen, G. von Ingersleben, and C.H. Chestnut III, "Problems with six point vertebral morphometry," *Proc. SPIE Medical Imaging 1998: Image Processing,* Kenneth M. Hanson; Ed., **3338**, pp. 107-117, 1998.
8. L.R. Hedlund and J.C. Gallagher, "Vertebral morphometry in diagnosis of spinal fractures," *Bone and Mineral*, **5**, pp. 59-67, 1988.
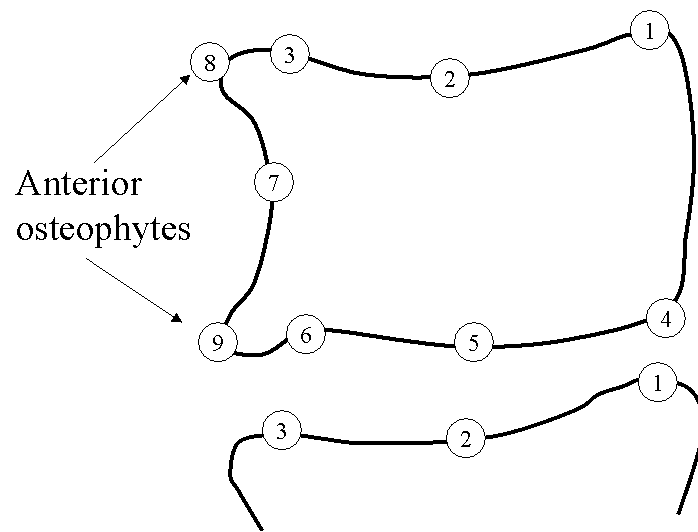
Figure 1. The nine radiologist-marked boundary points used to index the vertebrae
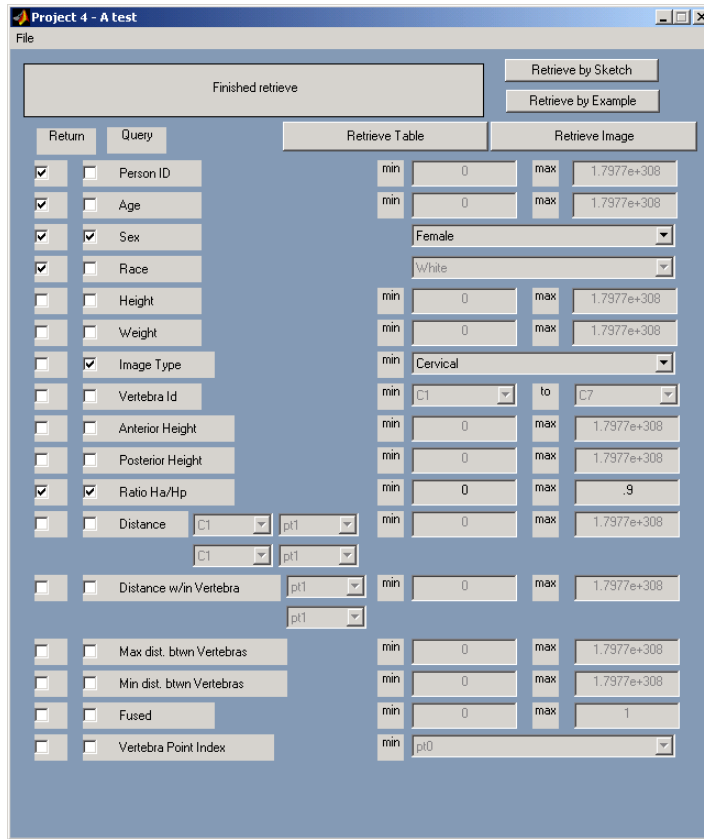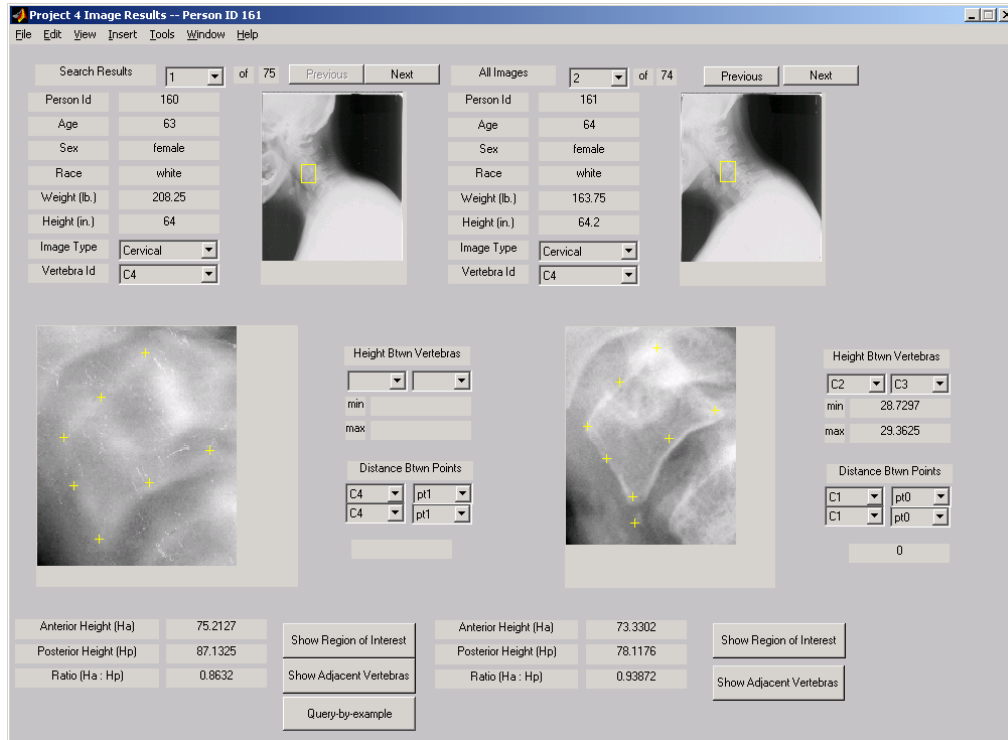
Figure 2. Query screen for pure text queries

Figure 3. Results of a pure text query. Shape similarity queries can be done on these results to do a hybrid text/shape query.
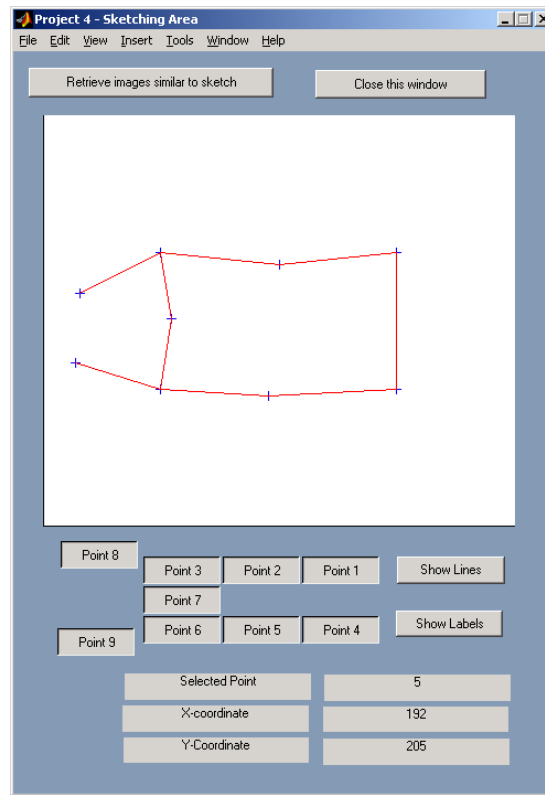
Figure 4. The example drawn shows large anterior osteophytees pointed toward the anterior face.
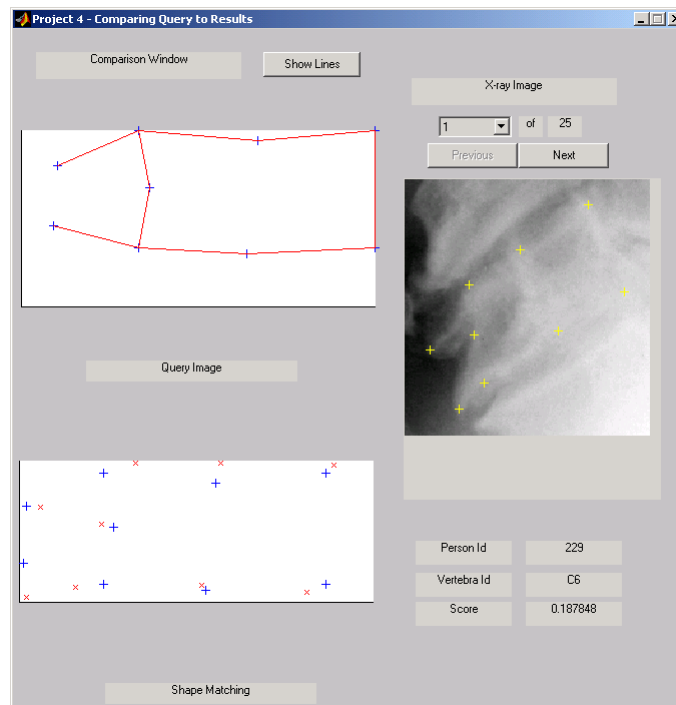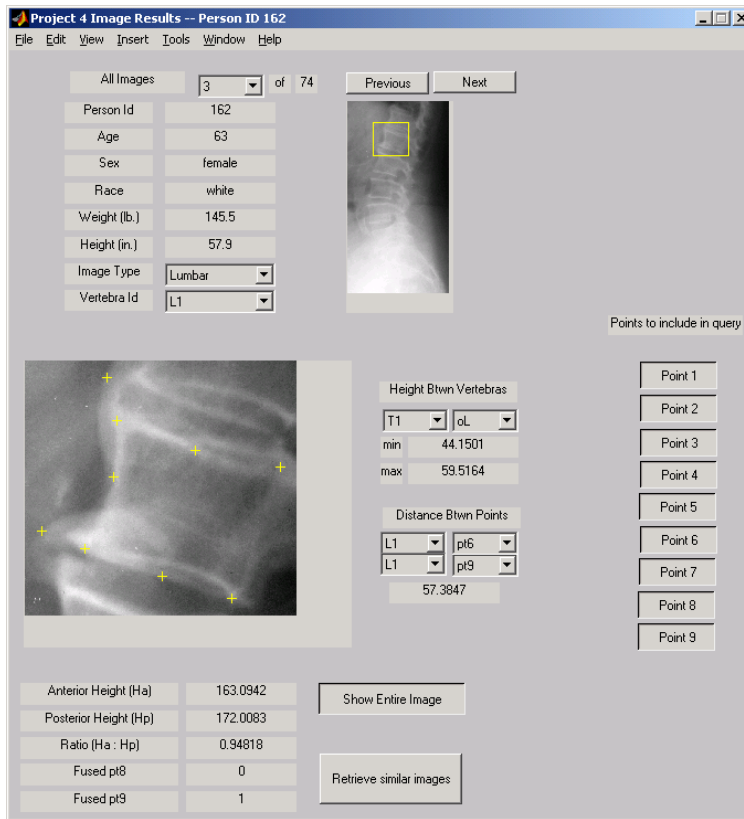

Figure 5. Results of query-by-sketch.

Figure 6. Query-by-example. The example image is a
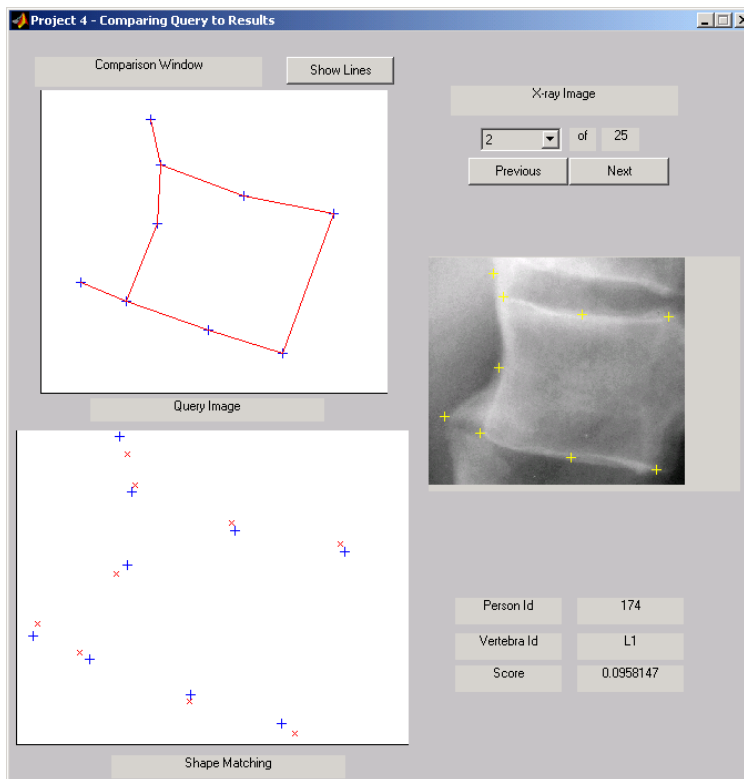lumbar vertebra with two prominent anterior osteophytes.

Figure 7. Results of query-by-example.