

# Recognition of Chest Radiograph Orientation for Picture Archiving and Communications Systems Display Using Neural Networks

John M. Boone, Sadananda Seshagiri, and Robert M. Steiner

**A neural network classification scheme was developed that enables a picture archiving and communications system workstation to determine the correct orientation of posteroanterior or anteroposterior chest images. This technique permits thoracic images to be displayed conventionally when called up on the workstation, and therefore reduces the need for reorientation of the image by the observer. Feature data were extracted from 1,000 digitized chest radiographs and used to train a two-layer neural network designed to classify the image into one of the eight possible orientations for a posteroanterior chest image. Once trained, the neural network identified the correct image orientation in 888 of 1,000 images that had not previously been seen by the neural network. Of the 112 images that were incorrectly classified, 106 were mirror images of the correct orientation, whereas only 6 actually had the caudal-cranial axis aligned incorrectly. The causes for misalignment are discussed.**

*Copyright © 1992 by W.B. Saunders Company*

**O**VER THE PAST DECADE there has been much discussion in the literature concerning the eventual demise of film, with the promise that computerized images will become the norm in radiology departments. Towards that end, so-called picture archiving and communications systems (PACS) have been developed at several institutions to serve as demonstration models of the all-digital radiology departments of the future. When a department invests in a digital image retrieval system, it is often necessary to digitize a large number of films to provide on-line recall of previous examinations in the PACS. The large-scale digitization of radiographic film is often performed using bulk feeders; therefore, the orientation of the film as it is digitized is often unknown to the PACS. In departments with computed radiography systems that use photostimulable phosphors, the

resulting digital images often exist on the system in an unknown orientation.

This study was performed to investigate an automatic method for determining the proper orientation of chest radiographs so that they may be correctly displayed without requiring the radiologist to reorient the image. Chest radiographs were used in this study because they represent a large percentage of radiographic images. Furthermore, examining previous chest radiographs is particularly important in diagnosis, and therefore the chest radiograph is a likely candidate for bulk digitization when a department converts its film library to a digital archive system.

## MATERIALS AND METHODS

One thousand posteroanterior (PA) chest radiographs were randomly selected and digitized to  $1,684 \times 2,048$  by 12 bits (4,096 shades of gray) using a laser film digitizer (Dupont FD-2000; Dupont Co, Wilmington, DE). The case selection included men and women with both normal and abnormal chest radiological diagnoses. All nonportable adult chest radiographs performed over a period of several months were included in the study. These films were sequentially digitized and were representative of the case mix at our institution. The images were stored on a write-once read-many (WORM) drive, resident on a 486-based microcomputer system (Northgate, Eden Prairie, MN). The system used a high-resolution ( $1,280 \times 1,024 \times 16$ -bit) display board (Dome Imaging Systems, Maynard, MA) for display of the digital images.

The images were reduced in size to  $256 \times 220$  pixels by averaging to accommodate more rapid processing of the image data. This lower resolution reduced the data used in processing by a factor of approximately 16, and in a short preliminary study it was found that the human observer could accurately identify the chest radiographic orientation with this resolution. Each image was displayed, and the correct orientation of the image was determined by a human observer and conveyed to the computer using a mouse-based interface. PA chest radiographs have eight possible orientations as shown in Fig 1. Orientations 1 through 4 represent the chest image with the caudal-cranial direction facing north, east, south, and west, respectively. In orientations 5 through 8, the caudal-cranial axes are positioned in the same direction as in 1 through 4, but the left-right axis is reversed. To succinctly refer to these categories, we define four directional categories A through D, as indicated in Fig 1. The two orientations in directional categories A and C are mirror images of each other; ie, they differ in their

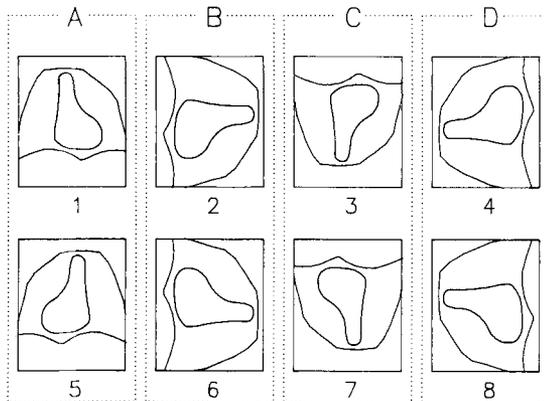
---

*From the Department of Radiology, Thomas Jefferson University, Philadelphia, PA.*

*Supported in part by National Cancer Institute grant R01-CA55372-01.*

*Address reprint requests to John M. Boone, PhD, 3390 Gibbon Bldg, Department of Radiology, Thomas Jefferson University, 111 S 11th St, Philadelphia, PA 19107.*

*Copyright © 1992 by W.B. Saunders Company  
0897-1889/92/0503-0007\$03.00/0*



**Fig 1. The eight possible orientations of a PA chest radiograph. Orientations 1 through 4 are 90° rotations of the image, and orientations 5 through 8 are the mirror images of 1 through 4. There are therefore two components of the orientation: direction, which refers to the caudal-cranial orientation, and mirror reversal.**

left-right orientation. Categories B and D differ in their up-and-down orientation.

To teach a computer to recognize the orientation of an image, certain features on the image need to be defined and examined by the computer. An image has a large number of features that can be defined, but the problem is to select only those features that will provide the necessary information for the task at hand. The features chosen for the present task are the vertical and horizontal projections of the image. The  $256 \times 220$  images were further reduced by averaging to yield  $32 \times 26$ -pixel images. These low-resolution images are "boxy" in appearance but readily reflect their proper orientation to the human observer. The vertical projections were calculated on the  $32 \times 26$ -pixel images by summing the counts in each of the pixel columns, generating 26 sums. The horizontal projections were generated by summing pixel rows, resulting in 32 more sums. All pixels within 5 pixels of the top or bottom edge of the image were excluded from the columnar sums, and all pixels within 5 pixels of the left or right edge were excluded from the row sums. In addition, four  $6 \times 6$ -pixel (on the  $32 \times 26$  image) regions of interest (ROIs) were summed. These four ROIs were the four quadrants of a  $12 \times 12$ -pixel region centered on the image. The motivation for these values was to provide additional information concerning the position of the mediastinum. The two projection profiles were combined into a single 58-element profile, which for each image was normalized to unity. The 4 ROI values, separately normalized to unity, were appended to this, resulting in an input data set of 62 values for each image.

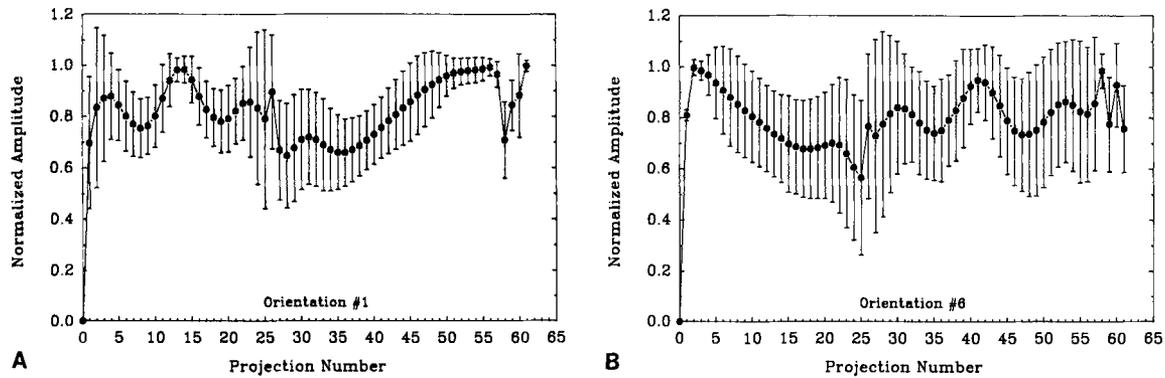
Of the 1,000 PA chest images used, 197 were oriented crosswise (eg, with the width of the patient along the long axis of the film) and 803 were lengthwise. To assure that each of the eight orientations was represented in reasonable numbers, the images were randomly but approximately uniformly distributed into the appropriate categories using a computerized randomization scheme. The original orientation was converted to the desired orientation using vertical and/or horizontal flipping of the images in com-

puter memory. The 197 crosswise radiographs were distributed among the even orientations (2, 4, 6, and 8), and the remaining 803 lengthwise images were distributed among the four odd-orientation categories.

Neural networks<sup>1</sup> are a form of artificial intelligence that learn to recognize patterns and differ substantially from conventional algorithm-based or rule-based artificial intelligence techniques. The training of a neural network is similar in concept to the training of a radiology resident in that both learn by example and assimilate knowledge from a large number of cases. A trained neural network is simply software that, when given a set of input data, responds with a set of output data. The input data are loaded into the input nodes of the neural network, and the output data are then retrieved at the output nodes of the network. Boone et al<sup>1</sup> give a simple explanation of neural network architecture. The neural network used the 62 projection values as input nodes. The output layer included eight nodes, one for each orientation. One hidden layer was used, with 10 hidden nodes. The neural network software was developed in our laboratory and used error back-propagation based on the generalized delta rule.<sup>2</sup> The 1,000 images were divided into two sets; set A included the first 500 images, and set B, the last 500 images. Set A was used in training the neural network, and the trained network was tested on set B. Conversely, set B was used in training the network, and set A was used in testing. The results from set A and set B were combined (training sets and testing sets separately) for analysis. Set A was trained using 1,546 iterations to a remaining error of 0.5%, and set B was trained using 2,957 iterations to a remaining error of 1.25%. The iterations represent the number of times the learning algorithm cycled through the data set. Each iteration results in an incremental improvement in the training, up to a point at which no further improvements are achieved with further iterations. The total training time was 3.44 hours on a 33-MHz 486 computer. The time required to process one image through the neural network using the neural network is imperceptible and was calculated to be approximately 1 millisecond. The training results reflect the capability of the input data to uniquely identify the output data and can in some cases reflect the appropriateness of the features that were selected. It is possible, however, that the neural network can learn to memorize the training set data and not truly learn to recognize the nuances in the input data that are useful in classification. The most valuable measure of neural network performance is the test set results, which reflect the ability of the neural network to generalize its acquired knowledge to data not previously seen by the network.

## RESULTS

The vertical and horizontal profiles and ROI values, which were used as input data to the neural network, were analyzed for set A. The mean profiles are shown in Fig 2 for two orientations. The mean value is shown at each point in the profile, along with error bars corresponding to  $\pm 2\sigma$ . The left portion (projections 1 to 26) of Fig 2A shows the vertical



**Fig 2.** The input data averaged over the images, with error bars corresponding to  $\pm 2\sigma$ . The  $2,000 \times 2,000$ -pixel images were reduced to  $26$  (wide)  $\times$   $32$  (high) pixels for feature analysis. The first  $26$  projection numbers correspond to the normalized column sums of the image, and the next  $32$  projection values ( $27$  through  $58$ ) are the normalized row sums of the image. The last four projection numbers represent the integrated pixel data in four regions in the center of the image. (A) The input data for orientation  $1$ , clearly showing the mediastinum as a bump centered around projection  $13$ . The positive slope in the data between projections  $27$  through  $58$  reflects the increase in film brightness (integrated horizontally) going down the image from the shoulders to the diaphragm. (B) The input data for orientation  $6$ , showing similar features as in A but rearranged because of the change in orientation.

projections of orientation  $1$ , and the profile clearly shows the peak corresponding to the mediastinum. Projections  $27$  through  $58$  represent the horizontal projection, and the gradual rise from left to right shows the projections going from the patient's shoulders, through the mediastinum, and to the diaphragm. The last four values in the data (projections  $59$  to  $62$ ) are the normalized sums of the four central region of interests. Figure 2B shows another profile, for orientation  $6$ . It is clear from these profiles that information concerning some aspects of the orientation of the image is present in the data.

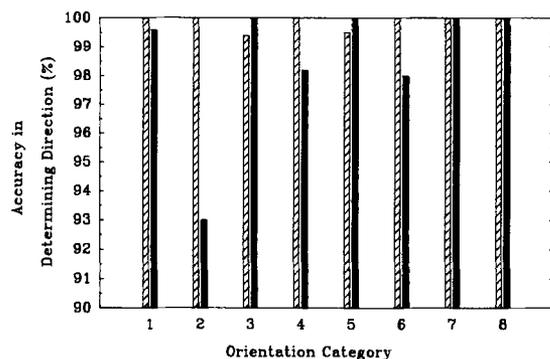
The ability of the neural network to determine the correct directional category was evaluated separately from the mirror-image results, and the results are shown in Fig 3. When applied to the training sets, the neural network failed to recognize the correction direction in only  $2$  of the  $1,000$  images. The test set classification of direction was also excellent, failing in only  $6$  of the  $1,000$  images ( $99.4\%$  accuracy).

In those images in which the directional category was correctly identified, the neural network distinguished the correct left/right position from the mirror image in  $98.1\%$  of images in the training set and in  $89.3\%$  of images in the test set. The accuracy for each of the eight orientations is shown in Fig 4 for both the training and testing data. Of the  $1,000$  images assessed in testing,  $888$  were perfectly oriented by the neural network. Of the  $112$  that were

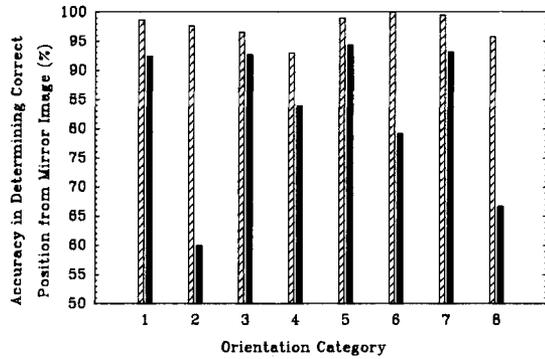
incorrect,  $106$  were classified by the neural network as the mirror image of the correct orientation. Of those  $106$  images that were mistaken for their mirror images,  $51$  of these were mistaken for crosswise images and  $55$  for lengthwise images. Because there were five times as many lengthwise images in the data set than crosswise, the error rates for the crosswise images ( $25.9\%$ ) is greater than that for lengthwise images ( $6.8\%$ ).

## DISCUSSION

The neural network technique was successful in recognizing the orientation of radiographs in a large majority of cases. The images were



**Fig 3.** The accuracy in determining the direction of the image, reflecting the ability of the technique to orient the clavicle up and the diaphragm down when the image is displayed. The direction was determined incorrectly in only  $6$  of  $1000$  images in the test set, and  $3$  of these were in orientation category  $2$ . ▨, Training results; ■, testing results.



**Fig 4.** The accuracy in determining the correct orientation from the mirror-image orientation for those images in which direction was identified correctly (eg, 99.4% of the images). The neural network had an 89.3% chance of correctly identifying the correct image from its mirror image. ▨, Training results; ■, testing results.

displayed with the caudal-cranial direction correct 99.4% of the time. From a computer-processing standpoint, flipping the image horizontally is a little faster than having to make 90° changes in the direction of the image. This technique should be useful in saving a physician time at a PACS workstation when chest radiographs are called up onto the screen. Although the time saved on a per-image basis is short, to frequent users of such systems these savings can be substantial over the lifetime of the PACS.

The crosswise images had a 25.9% chance of being mistaken for their mirror images, as opposed to a 6.9% chance of a similar mistake

for lengthwise images. This fact is probably a consequence of the following: (1) there were fewer crosswise images in the training sets, and therefore the neural network had fewer examples in which to learn; and (2) the mediastinum has a greater leeway in terms of its position on the image in crosswise images, because the left-to-right (mirror) direction is longer in this orientation.

The image data base used in this study was representative of the caseload at a large urban hospital, and there were many images interpreted as significantly abnormal in these 1,000 cases. For example, for cases in which pulmonary edema was present, the mediastinum was difficult to identify, and therefore correct classification was impaired.

**CONCLUSIONS**

Various features were extracted from the digital chest radiographs and were used in conjunction with a neural network to accurately predict the correct orientation of the chest radiographs. This procedure should substantially reduce the frequency at which chest radiographs need to be reoriented after being called up for display on a PACS workstation. Therefore, this relatively simple technique may save physicians time and frustration, especially when considered over the useful life of a PACS imaging workstation.

**REFERENCES**

1. Boone JM, Gross GW, Greco-Hunt V: Neural networks in radiologic diagnosis: Introduction and illustration. *Invest Radiol* 25:1012-1016, 1990
2. Rumelhart DE, McClelland JL, PDP Research Group: *Parallel Distributed Processing*, vols 1 and 2. Cambridge, MA, Massachusetts Institute of Technology, 1988