COMPUTERS AND BIOMEDICAL RESEARCH 25, 569–585 (1992)

A Neural Network Architecture for Understanding Discrete Three-Dimensional Scenes in Medical Imaging

G. COPPINI,* R. POLI,† M. RUCCI,† AND G. VALLI†

*C.N.R. Institute of Clinical Physiology, Pisa, Italy; and †Department of Electronic Engineering, University of Florence, Florence, Italy

Received October 9, 1991

Magnetic resonance and computed tomography produce sets of tomograms which are termed discrete 3D scenes. Usually, discrete 3D scenes are analyzed in two dimensions by observing each tomogram on a screen so that the three-dimensional information contained in the scene can be recovered only partially and qualitatively. The three-dimensional reconstruction of the shape of biological structures from discrete 3D scenes would allow a complete and quantitative recovery of the available information, but this task has proved hard for conventional processing techniques. In this paper we present a system architecture based on neural networks for the fully automated segmentation and recognition of structures of interest in discrete 3D scenes. The system includes a retina and two main processing modules, an Attention-Focuser System and a Region-Finder System, which have been implemented by using feed-forward nets trained with the back-propagation algorithm. This architecture has been tested on computer-simulated structures and has been applied to the reconstruction of the spinal cord and the brain from sets of tomograms. © 1992 Academic Press, Inc.

1. INTRODUCTION

Magnetic Resonance (MR) and Computed Tomography (CT) are among the most widely applied digital techniques in medical imaging. They produce a set of tomograms representing parallel sections of the body. According to Udupa (*1*), the set of 2D digital images is termed *discrete 3D scene* and can be seen as a 3D array of items, called voxels, whose values depend on the local characteristics of biological tissues.

Usually, discrete 3D scenes are analyzed in two dimensions by observing the set of tomograms on a screen. In this way the observer can recover the three-dimensional information contained in the scene partially and qualitatively. This is the basic reason why in the last few years methods have been developed to analyze discrete 3D scenes quantitatively (*2–6*).

In this view, three-dimensional reconstruction of the shape of the imaged structures is one of the most attractive goals. In fact, operations like quantitative measurements, surgery planning, and radiation therapy can take advantage of

0010-4809/92 $5.00

such a reconstruction (*7–12*). Furthermore, shape reconstruction methods can be usefully applied to imaging techniques different from MR and CT, e.g., histological imaging (*13–16*). In addition, a 3D description of the shape of biological structures can be utilized for scene-understanding tasks such as recognition, morphological analysis, and clinical decision-making. However, given the variability and complexity of anatomical structures as well as the effects of image noise, fully automatic systems for reconstructing biological shapes have not been developed.

According to the most widely accepted computer vision theories (*17*), segmentation and understanding of discrete 3D scenes can be reached following several strategies which lie between the so-called bottom-up (or data-driven) and top-down (or model-driven) approaches.

In general, plain bottom-up strategies have not shown adequate sensitivity and specificity. As emphasized by Marr (*18*), in the case of segmentation this is mainly due to the lack of an exact formulation of the task. In fact, regions with high semantic importance do not always have particular visual distinctions.

On the other hand, with a fully top-down strategy this problem does not exist, as the considered visual process is driven by internal models producing expectations on the input data. However, the accuracy and flexibility of internal models play a crucial role.

Due to their typical properties, Artificial Neural Networks (ANN) (*19, 20*) can automatically integrate bottom-up and top-down paradigms so as to overcome the problems described above. In fact, the output of a neural network is dependent both on the input data and on the a priori expectations that have been stored in the network connections during the training phase. Thus, using ANNs there is no need for an explicit description of the imaged object because such a description is implicitly included in the net in a distributed form.

In this paper, we are proposing a neural architecture for segmenting and recognizing biological structures in discrete 3D scenes as produced by MR and CT scanners.

The basic design requirements include the recognition of the part of the human body being imaged (such as the head or thorax), which we term *anatomical district*, the identification of the organs of interest in the anatomical district, and the 3D reconstruction of their shape. In our opinion, to achieve these goals the use of a multiresolution approach and of attention-focusing mechanisms is advisable. Both are of biological inspiration and are suitable for an ANN implementation; they also allow a drastic reduction of computation load.

A few of the existing segmentation techniques process the scene as a single 3D image, thus taking advantage of correlation among the data along the scanning axis. On the other hand, these techniques usually require large amounts of memory and a noteworthy computation power. In addition, they cannot be applied to sparse or unevenly spaced sets of slices (*21*). For this reason, we chose to process discrete 3D scenes in a slice-wise way, restricting the use of the information of the third dimension to the evaluation of the *characteristic view* of each tomogram. The characteristic view is a member of a minimal set

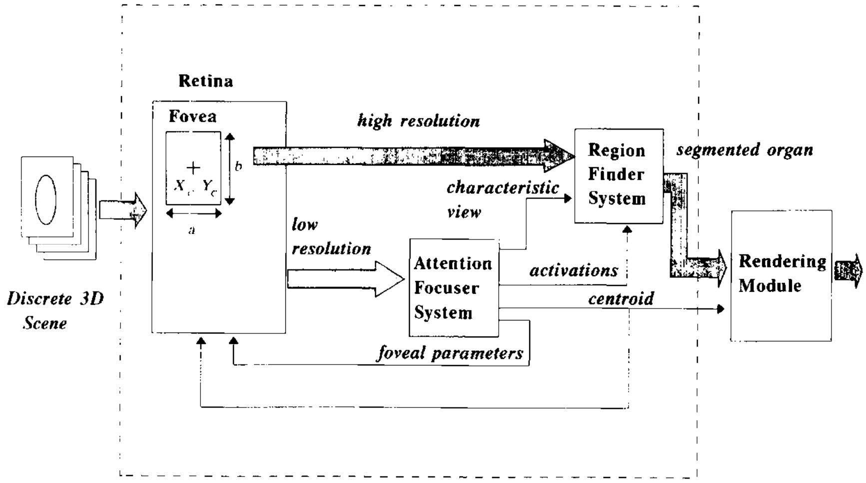

FIG. 1. System architecture.

of prototypal images (which differ from each other in the number or the shape of the imaged organs) to which the analyzed image is most similar.

2. System Architecture

The system architecture we developed is schematized in Fig. 1. Two main blocks can be singled out: the one which is enclosed in the dashed box builds a representation of the input scene, and the other improves this representation to produce an adequate display. In this paper we focus our attention on the first block. It includes an input section, which we term retina, an Attention-Focuser System (AFS), and a Region-Finder System (RFS).

As previously mentioned, the system processes a slice at a time. When the first image of the scene is presented to the system, the retina decimates and smoothes it. A region of the retina, called the fovea, performs higher resolution decimation and smoothing on a reduced area of the input image, termed the Region of Interest (ROI). The position and the dimensions of the fovea are controlled by signals coming from the AFS, as described later.

Then, the low-resolution image provided by the retina is processed by AFS. At first it identifies the imaged anatomical district d ($d = 1, \ldots, D$, where D is the number of considered districts) and evaluates the code z_d of the characteristic view for the input tomogram ($z \in \{1, \ldots, V_d\}$, where V_d is the number of characteristic views for district d). Then, AFS focuses on an organ of interest at a time. For each organ it provides the coordinates X_c, Y_c of the centroid of the organ and other fovea parameters such as the dimensions (a, b). Once these data are known, the fovea processes the ROI. Its output is a representation of the organ of interest which is insensitive to translations. The fovea output and the information about the district and the organ being analyzed are sent to the

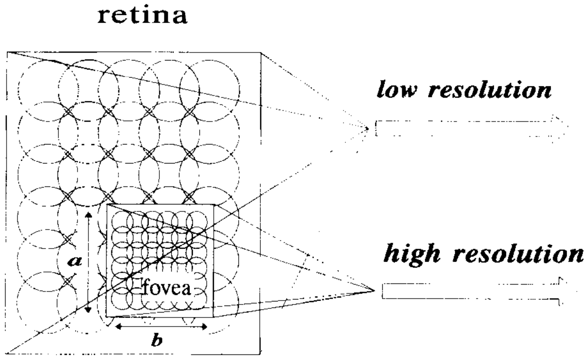

FIG. 2. The retina.

RFS, which extracts the pixels belonging to the organ. The resulting bit map and the centroid coordinates are stored.

After an organ bit map has been built, AFS focuses on another organ until all the organs in the slice have been segmented. After the first slice has been completely analyzed, the whole process is repeated for the other images of the scene.

Summarizing, the output representation of the discrete 3D scene includes a description of each organ which consists of:

i) a generalized axis made up of the centroids provided by the AFS;

ii) a stack of bit maps supplied by the RFS, each bit map representing a cross section of the organ.

3. A Neural Network Implementation

In this section we describe an implementation of the previously outlined architecture which is mainly based on feed-forward nets trained with the back-propagation algorithm (*22*).

The retina can be seen as a neural network made up of an input layer, including as many neurons as image pixels, and an output layer which usually includes 16×16 linear neurons. Like in biological retinas, the connections between the input and the output neurons are local and their strength has a bidimensional Gaussian profile (typically we used $\sigma = 4.5$) centered on an output neuron. These Gaussian profiles cover partially overlapping areas of the image shown as circles in Fig. 2. Actually, the process performed by retina can also be implemented by filtering the input image with a Gaussian kernel and then resampling the filtered image. (This is computationally efficient, especially for software implementations on sequential computers.) The fovea has a similar structure, though with different parameters (see Fig. 2).

As shown in Fig. 3, AFS is made up of several modules: a District-Selection Module (DSM), a Level-Evaluation Module for each considered anatomical

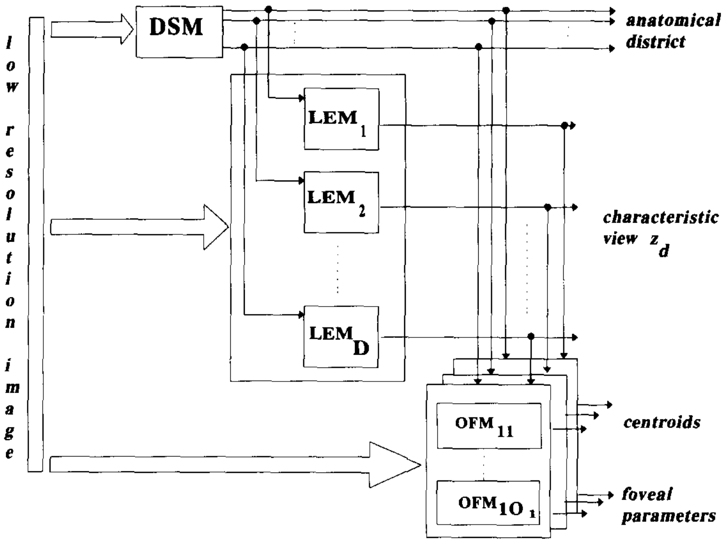

FIG. 3. The modules of the Attention-Focusing System (AFS). This includes a District-Selection Module (DSM), a Level-Evaluation Module for the dth anatomical district (LEM_d), and an Organ-Focuser Module for the oth organ in the dth district (OFM_{do}).

district d (LEM_d), and an Organ-Focuser Module for each organ of interest o in the selected district d (OFM_{do}).

On the basis of the observed slice, DSM identifies the imaged anatomical district and produces a set of activation signals which control the operation of LEMs and OFMs as well as of the modules of RFS. DSM is a four-layer fully connected network including 32 + 32 hidden units. The number of its input units is equal to that of the output neurons of the retina, while the output units are D.

Once DSM recognizes district d, it activates LEM_d, which estimates the code z_d of the characteristic view of the input image. LEMs (see Fig. 4) work as follows. First, the centroid of the whole section is estimated by a focusing module similar to OFM which is described below. The centroid computed by the focusing module is taken as the center of a polar reference (r, θ) and each pixel is spatially relocated into ($\ln(r + 1)$, θ) (logarithmic–polar coordinate transformation). This transformation can be seen as a nonisotropic scale change

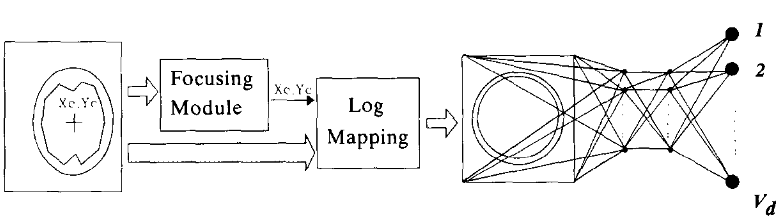

FIG. 4. Functional scheme of a LEM.

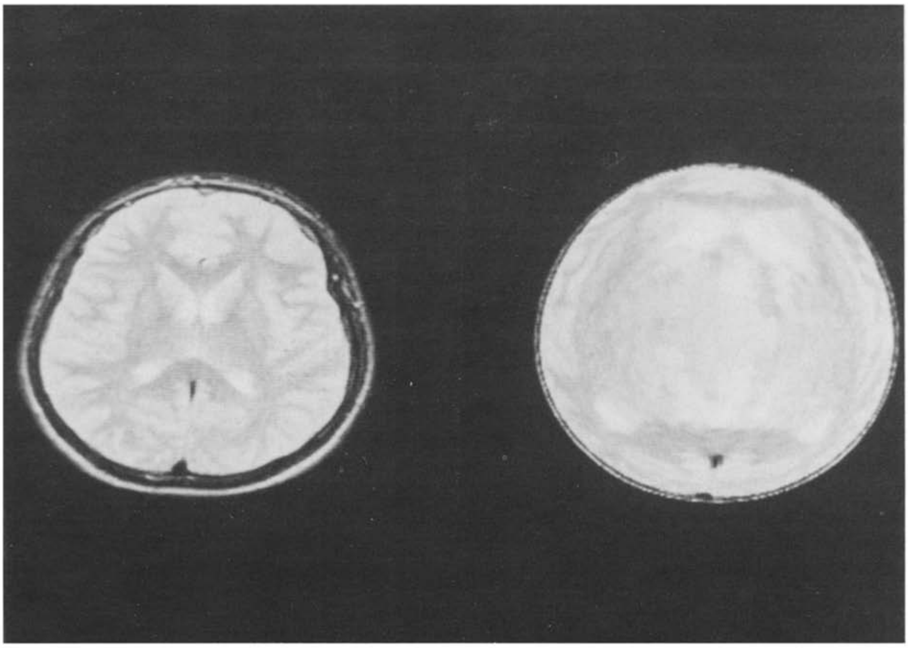

FIG. 5. Example of logarithmic–polar coordinate transformation: original image on the left and transformed image on the right. For display convenience, transformation has been carried out at the highest resolution on a 256 × 256 matrix (in the actual system operation this is performed at the lowest resolution).

of the image which stretches the image around the centroid and shrinks it elsewhere. Then, a four-layer fully connected network estimates z_d. This network is made up of 16 × 16 input units, 32 + 32 hidden units, and V_d output units. The log–polar mapping reduces the importance of scale changes so that this neural network operates mainly on the basis of the morphology of the imaged structure rather than on the basis of its dimensions. An example of logarithmic–polar transformation is shown in Fig. 5.[1]

Once the characteristic view is known, OFM_{do} ($o = 1, \ldots, O_d$, where O_d is the number of organs in district d) are activated sequentially. Each OFM estimates the centroid of the related organ. OFM also gives the dimensions

[1] The reader should note that, in general, logarithmic transformations such as that performed by LEM depend on the center position. However, in our case this problem does not influence LEM performances. In fact: (a) errors made in centroid evaluation are small (the maximum error has been 5 pixels for 256 × 256 images), and (b) characteristic views are quite different from each other so that they can be reliably discriminated even in the presence of distortions introduced by inexact logarithmic transformation.

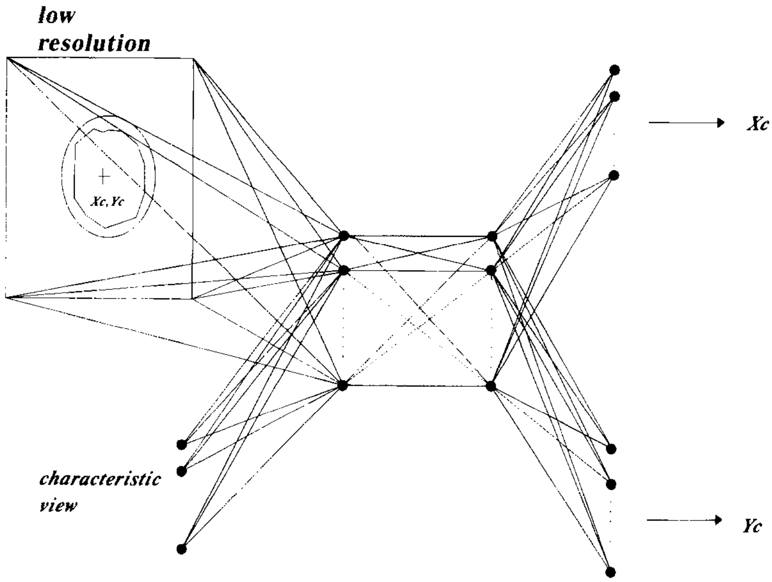

FIG. 6. Scheme of an OFM network.

(a_{do}, b_{do}) of the foveal region and σ_{do} of the related Gaussian kernel: they are predefined according to the size of the organ of interest. Subsequently, the fovea can be positioned on the ROI.

As shown in Fig. 6, each OFM includes a four-layer fully connected net. It processes two different kinds of information. In fact, the network has (a) 16×16 input units fed by the retina and (b) V_d input units from LEM_d. In addition, the net includes 32 + 32 hidden units and 32 output units. These are grouped in two sets of 16 units (one set for each reference axis) and give, with a sparse coding (see (*23*)), the coordinates of the centroid.

The Region-Finder System (see Fig. 7) includes a set of Segmenter Modules (SM_{do}, one for each organ o of each considered district d) which classify the pixels belonging to the related organs. Each SM is activated after the related OFM has completed its task. An SM consists of a five-layer neural net. The network inputs are (a) the gray levels of the ROI pixels contained in a square region (called mask), (b) the position of the mask in the ROI (sparsely coded by the coordinates of its central pixel), and (c) the characteristic view of the tomogram as evaluated by AFS. As shown in Fig. 8, the different kinds of information (gray levels and position) are processed independently by two different fully connected subnetworks joined in a common layer. The activation of the output unit indicates whether this pixel belongs to the considered organ.

Since the ROI dimensions depend upon the considered organ, the number of units which make up the position-processing subnetwork can be different for each SM. A typical network topology includes 16 input units for each coordinate

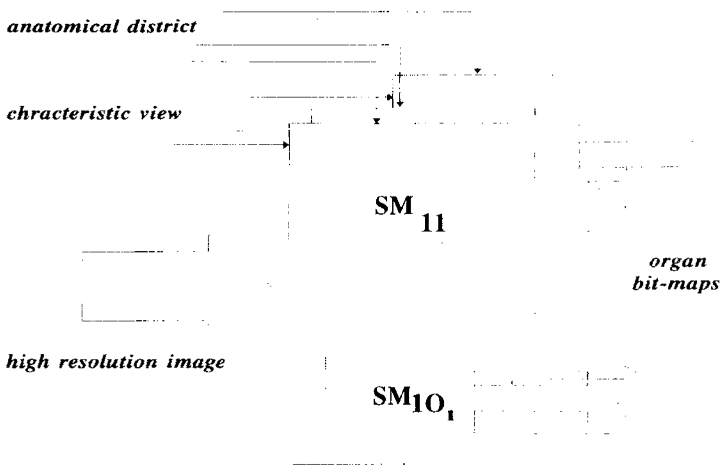

FIG. 7. Scheme of an RFS.

axis, V_d input units for the code z_d of the characteristic view, and 16 + 16 hidden units. The gray-level-processing subnetwork has a fixed number of units: 9 × 9 input units, 10 units in the first hidden layer, and 5 units in the third one. By using a suitable threshold, the output of SM networks is converted into binary values, thus giving a bit map representation of the organ section.

We point out that, since the output of each SM results from a combination of

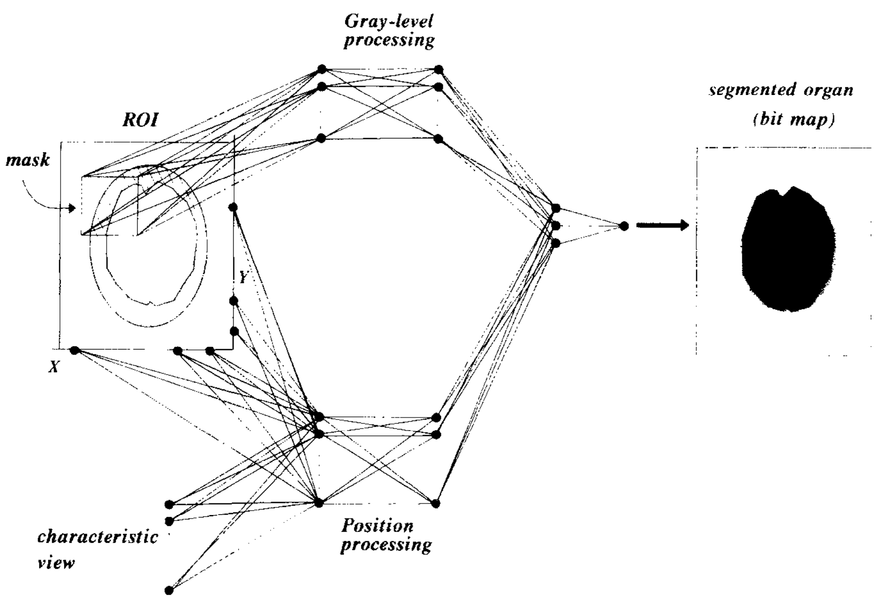

FIG. 8. Scheme of an RFS network. (x, y) identifies the position of the mask.

TABLE 1

PHANTOM COMPONENTS

Ellipsoid	Center coordinates	Major axis	Minor axis	Zeta axis	Rotation angle	Gray level
A	127 ± 5, 127 ± 5, 0 ± 5	117 ± 3	88 ± 3	90 ± 5	0	100
B	127 ± 5, 130 ± 5, 0 ± 5	111 ± 3	84 ± 3	85 ± 5	0	150
C	156 ± 5, 127 ± 5, 0 ± 5	40 ± 3	14 ± 3	30 ± 5	1.88	190
D	99 ± 5, 127 ± 5, 0 ± 5	52 ± 3	20 ± 3	30 ± 5	1.25	190
E	127 ± 5, 88 ± 5, 0 ± 5	32 ± 3	27 ± 3	50 ± 5	0	140
F	127 ± 5, 115 ± 5, 0 ± 5	6 ± 3	6 ± 3	6 ± 5	0	200
G	127 ± 5, 140 ± 5, 0 ± 5	6 ± 3	6 ± 3	6 ± 5	0	120
H	117 ± 5, 205 ± 5, 0 ± 5	6 ± 1	6 ± 1	6 ± 3	0	180
I	127 ± 5, 205 ± 5, 0 ± 5	3 ± 1	3 ± 1	3 ± 3	0	180
J	135 ± 5, 205 ± 5, 0 ± 5	6 ± 1	3 ± 1	3 ± 3	1.57	180

gray-level data with the position of the pixels being classified, SM operation can be seen as an optimal nonlinear, space-varying filtering. It should be noted that the evaluation of the characteristic view of each slice made by LEMs greatly increases the performance of SMs. In fact, if these data were not available, the knowledge stored in SMs would be a mixture of the characteristic views of the analyzed district. On the contrary, the availability of z_d has allowed these nets to learn and then use specific knowledge about each characteristic view.

Once all the slices of the scene have been analyzed, the representations of the organs of interest are sent to a 3D rendering module. When needed it improves these representations with an interpolation process which computes a set of contours which fill in the gaps between adjacent slices. At present, the dynamic–elastic interpolation algorithm (*5*) is implemented. After the interpolation phase, the organs are visualized on a computer display with shading techniques for interactive analysis.

4. Results and Conclusions

The validity of the described approach was at first tested by means of simulation techniques. The use of a geometrical phantom has permitted a quantitative analysis of the performances of the overall system. In particular, the output of the segmentation modules was compared with the "ideal" segmentation given by the mathematical description of the test objects. Although simulation results cannot directly compare with those obtained from real discrete 3D scenes, nevertheless they can be used as a reference for understanding the system operation.

The employed test object is a 3D extension of the Shepp and Logan "head phantom" (*24*). The 10 ellipses which compose the original phantom were converted into 10 ellipsoids described in Table 1. To introduce shape variability, uniform deviates (whose standard deviations are given in the table) were employed to compute the actual ellipsoid parameters.

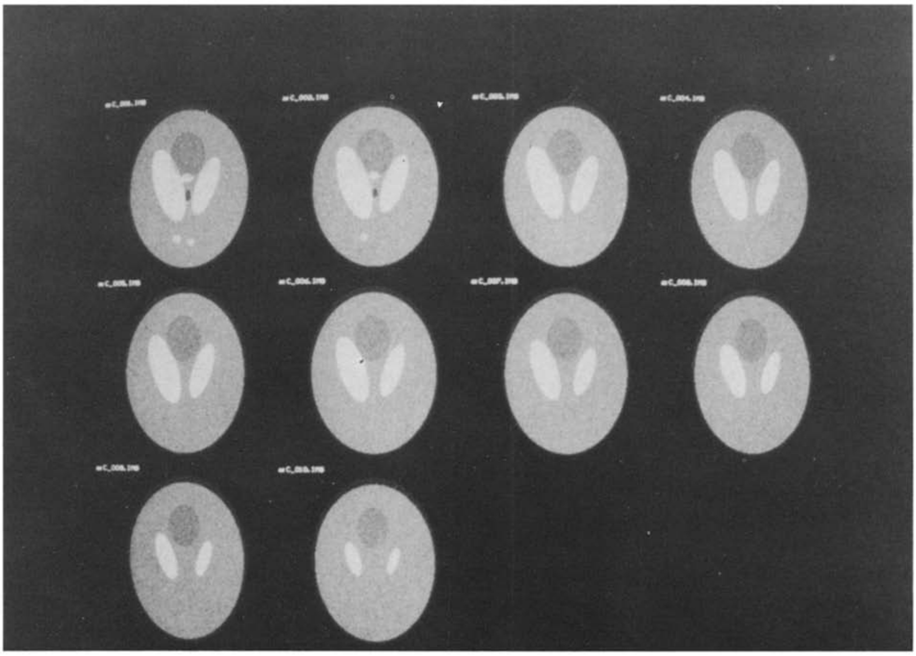

FIG. 9. Tomograms of a simulated discrete 3D scene.

Noise-free simulated discrete 3D scenes were obtained by scanning the phantom. The final synthetic scenes were built by adding a zero-mean Gaussian white noise to each simulated section. In Fig. 9, 10 sections of a simulated object are displayed.

In the following, the segmentation and reconstruction of ellipsoid E of the phantom are considered.

The neural networks were trained by using nine test sequences, each consisting of 10 evenly spaced slices. For each section the standard deviation of the superimposed noise was 10 gray levels. In this case, a set of six characteristic views was adopted.

A further set of simulated objects was generated by modifying the phantom parameters. By means of this set the system performance was evaluated. The rates of correctly and wrongly assigned pixels (CA and WA) were computed. In addition, the rates of correct classification of pixels lying inside (true positive, TP) and outside (true negative, TN) ellipsoid E were computed.

The obtained results are shown in Table 2. As reported in the first two columns, about 99% of the pixels are correctly classified. If we neglect the last row for the moment, we can say that CA has a maximum for slice 6. The values shown in the last two columns would suggest that, at least under these experimental conditions, the system is proficient in recognizing the pixels that do not belong to the considered structure. In fact, for six levels of the test scene

TABLE 2

SYSTEM PERFORMANCE ON A TEST DISCRETE 3D SCENE

Slice no.	CA	WA	TP	TN
1	0.985	0.015	0.921	1
2	0.986	0.014	0.925	1
3	0.988	0.012	0.935	1
4	0.992	0.008	0.952	1
5	0.992	0.008	0.950	1
6	0.997	0.003	0.985	0.999
7	0.995	0.005	0.992	0.996
8	0.994	0.006	1	0.994
9	0.991	0.009	1	0.991
10	1	0	1	1
Mean values	0.992	0.008	0.966	0.998

all pixels outside the ellipsoid were correctly classified (TN $=$ 1). Conversely, the classification of the pixels making up the ellipsoid improves while increasing the level of the section (TP $\approx$ 1). This may be due to the fact that in most sections the inner pixels of the ROIs belong to the ellipsoid.

The last row in Table 2 represents the results obtained from an image in which, due to the variability of the parameters of the phantom, ellipsoid E is absent and all pixels are correctly classified as being outside the structure of interest. This row has been included to show the behavior of the system in a situation in which data information prevails completely over the prior knowledge.

The noise-rejection capabilities of the system were evaluated on a series of images with increasing level of noise. All scenes were built with the same phantom parameters and the standard deviation of noise was varied from 10 to 80 gray levels, as illustrated in Fig. 10. In Fig. 11, the rates of wrong assignments, false positives, and false negatives are plotted against the noise standard deviation. These results refer to the segmentation of the sixth sectional image.

The prior knowledge embedded in the system is of fundamental importance when one is working with very noisy images. The curves in Fig. 11 indicate that the system performs pretty well even when the noise is strong. In particular, the figure shows that when the noise increases, the increment of wrongly classified pixels is mainly due to false positives, i.e., to the inclusion in the segmented structure of pixels which do not belong to the object. Even with high noise levels the system seldom rejects the pixels belonging to the structure of interest. Likely, this depends on the action of detailed knowledge about the typical shape of the structure being segmented. Figure 12 shows some examples of segmented phantom slices.

After these tests on phantom scenes, the system has been applied to real MR and CT images. First, the reconstruction of the spinal cord from CT sections

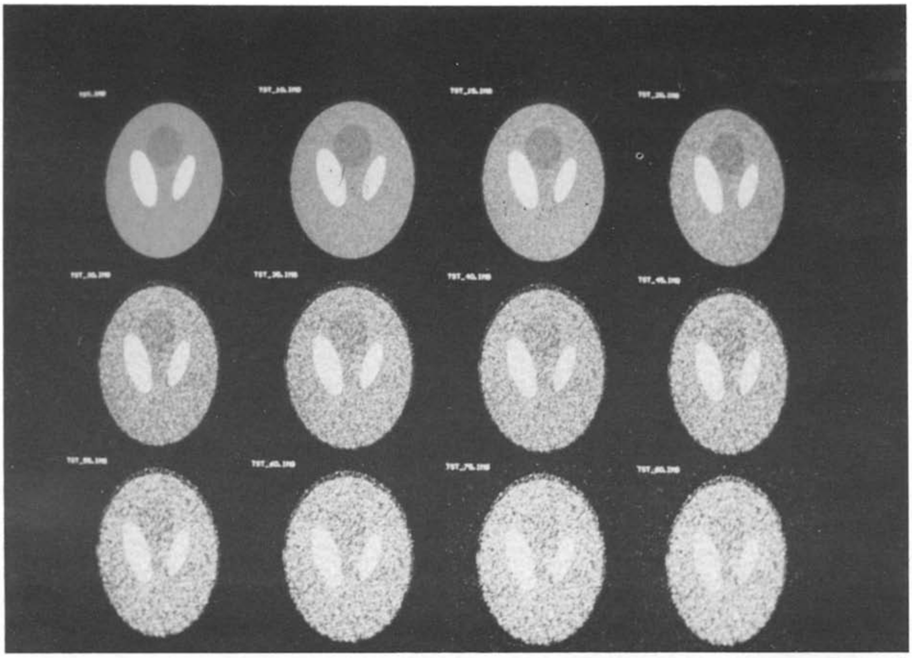

FIG. 10. A simulated tomogram with noise standard deviation varying from 10 to 80 gray levels.

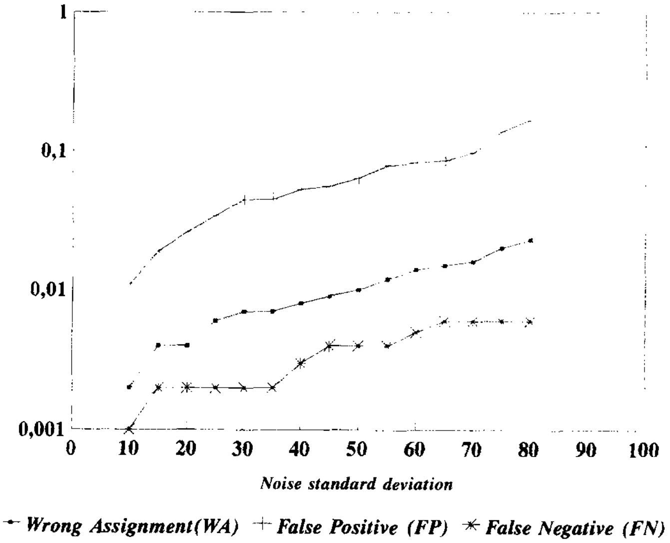

FIG. 11. System performance with varying noise levels for a simulated tomogram.

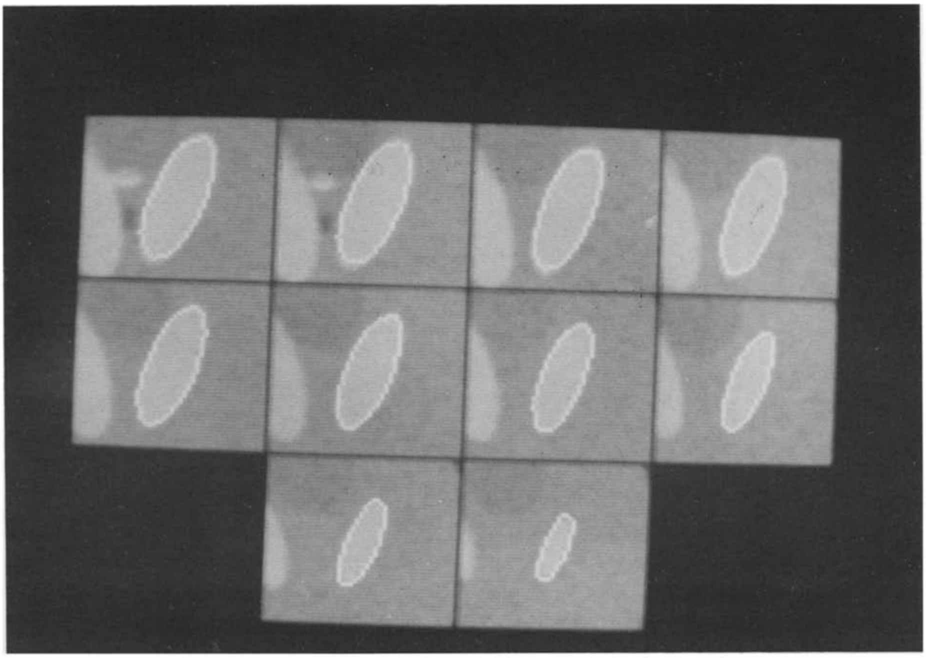

FIG. 12. Examples of segmented phantom images (not included in the training set). Only focused ROIs are shown.

was performed. The networks were trained with 40 thoracic slices from four patients. In these experiments a single characteristic view was used. The obtained ROIs were square-shaped and their side was 80 pixels wide. The performance of the system for the images of the four patients of the training set is shown in Table 3. The performances seem to be quite satisfactory. It is worth noting that all pixels belonging, in the trainer's opinion, to the considered structure have been correctly classified by the system.

Similar results were obtained when images not included in the training set

TABLE 3

SYSTEM PERFORMANCE FOR SPINAL CORD DISCRETE 3D SCENES (TRAINING SET)

Scene no.	CA	WA	TP	TN
1	0.993	0.007	1	0.990
2	0.993	0.007	1	0.992
3	0.997	0.003	1	0.995
4	0.989	0.011	1	0.986
Mean values	0.993	0.007	1	0.991

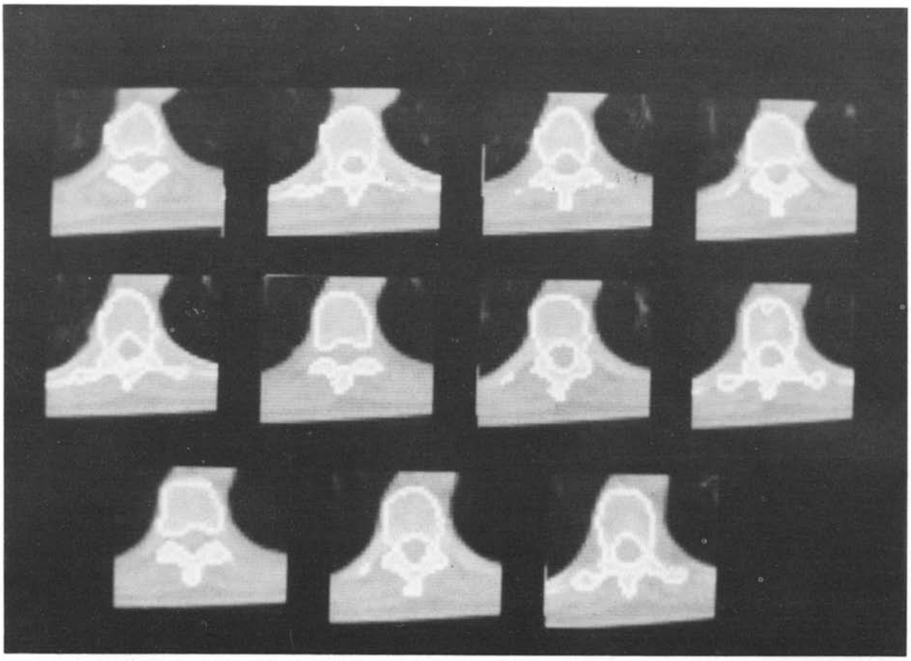

FIG. 13. Examples of segmented CT images of the spinal cord (not included in the training set). Only focused ROIs are shown.

were processed. An example of segmentation of a spinal cord is illustrated in Fig. 13.

The system was also used for processing the brain from MR scenes. In this case, the neural networks were trained on 36 images from three patients. For the head district, ROIs had the same dimensions as the original images, i.e., 256 × 256, and six characteristic views were singled out. The results obtained are listed in Table 4. Again the system seems to perform quite well. An example of segmentation of a brain is shown in Fig. 14.

Three-dimensional views of the spinal cord and the brain, as produced by the 3D rendering module, are illustrated in Fig. 15.

TABLE 4

SYSTEM PERFORMANCE FOR BRAIN DISCRETE 3D SCENES (TRAINING SET)

Scene no.	CA	WA	TP	TN
1	0.993	0.007	0.955	0.996
2	0.991	0.009	0.967	0.995
3	0.987	0.013	0.951	0.993
Mean values	0.990	0.009	0.957	0.994

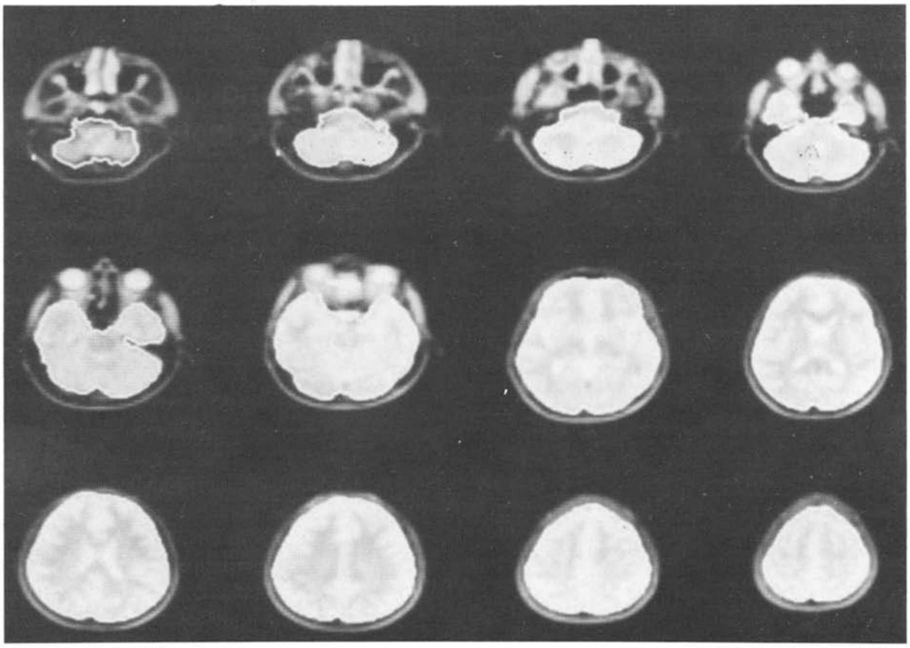

FIG. 14. Examples of segmented MR images of the brain (not included in the training set). Only focused ROIs are shown.

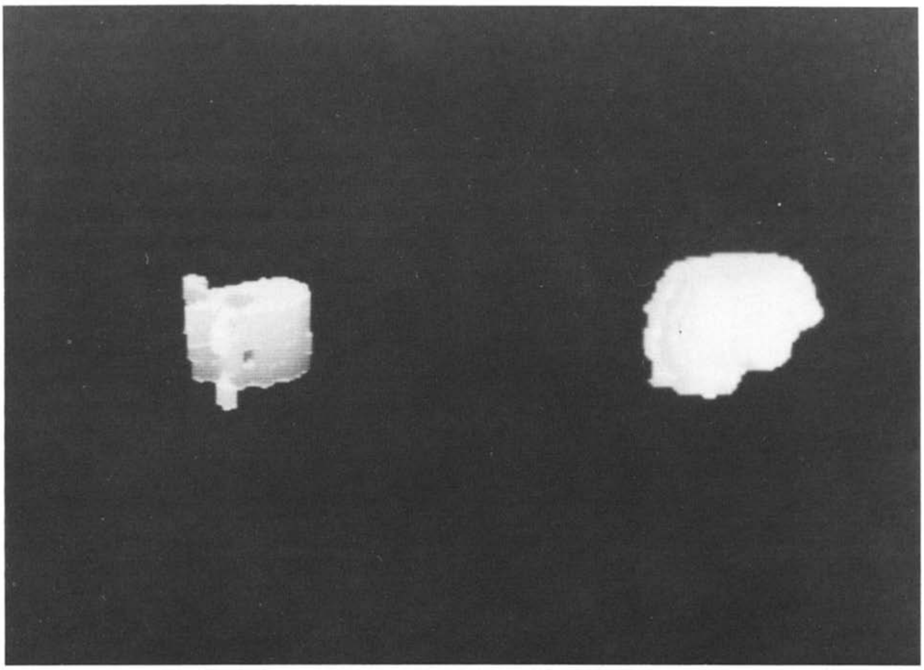

FIG. 15. Typical outputs of the rendering module.

In conclusion, our results support the concept that the integration of top-down and bottom-up processing performed by ANNs, coupled with the use of multiresolution analysis and attention-focusing mechanisms, allows the development of systems that are both reliable and automatic for complex tasks like segmenting and recognizing organs in discrete 3D scenes.

REFERENCES

1. UDUPA, J. K. Display of 3-D information in discrete 3-D scenes produced by computerized tomography. *Proc. IEEE* **71** (1983).
2. CHEN, L. S., AND SONTAG, M. R. Representation, display and manipulation of 3D digital scenes and their medical applications. *Comput. Vision Graphics Image Process.* **48,** 190 (1989).
3. FABER, T. L., AND STOKELY, E. M. Feature detection in 3D medical images using shape information. *IEEE Trans. Med. Imaging* **6,** 1 (1987).
4. GORDON, D., AND UDUPA, J. K. Fast surface surface tracking in three-dimensional binary images. *Comput. Vision Graphics Image Process.* **45,** 196 (1989).
5. LIN, W. C., LIANG, C. C., AND CHEN, C. T. Improvement on dynamic elastic interpolation technique for reconstructing 3-D objects from serial cross sections. *IEEE Trans. Med. Imaging* **9,** 1 (1990).
6. UDUPA, J. K., AND AJJANAGADDE, V. G. Boundary and object labeling in three-dimensional images. *Comput. Vision Graphics Image Process.* **51,** 355 (1990).
7. BOISSONAT, J. D. Shape reconstruction from planar sections. *Comput. Vision Graphics Image Process.* **44,** 1 (1988).
8. KENNEDY, D. N., AND NELSON, A. 3-D display from cross sectional tomographic images: An application to magnetic resonance imaging. *IEEE Trans. Med. Imaging* **6,** 2 (1987).
9. RICHARD, M. J., ALLARD, J., GHOSH, S. K., AND BOUGOUSS, M. Three-dimensional reconstruction of human limbs from tomographic views. *Comput. Biomed. Res.* **22,** 26 (1989).
10. ROBB, R. A., AND BARILLOT, C. Interactive display and analysis of 3-D medical images. *IEEE Trans. Med. Imaging* **8,** 3 (1989).
11. TATSUMI, H., TAKAOKI, E., OMURA, K., AND FUJITA, H. A new method for three-dimensional reconstruction from serial sections by computer graphics using "meta-balls": Reconstruction of "hepatoskeletal system" formed by Ito cells in the cod liver. *Comput. Biomed. Res.* **23,** 37 (1990).
12. BLOCH, P., AND UDUPA, J. K. Application of computerized tomography to radiation therapy and surgical planning. *Proc. IEEE* **71,** 3 (1983).
13. JOHNSON, E. M., AND CAPOWSKY, J. J. A system for the three-dimensional reconstruction of biological structures. *Comput. Biomed. Res.* **16,** 79 (1983).
14. PROTHERO, J. S., RIGGINS, M., LINDSAY, A., HARRIS, R., AND PROTHERO, J. Three-dimensional reconstruction from serial sections. III. Autoscan, a software package in Fortran for semiautomated photomicrography. *Comput. Biomed. Res.* **18,** 132 (1985).
15. PROTHERO, J. S., AND PROTHERO, J. W. Three-dimensional reconstruction from serial sections. IV. The reassembly problem. *Comput. Biomed. Res.* **19,** 361 (1986).
16. WONG, Y. M. M., THOMPSON, R. P., COBB, L., AND FITZHARRIS, T. P. Computer reconstruction of serial sections. *Comput. Biomed. Res.* **16,** 580 (1983).
17. BALLARD, D., AND BROWN, B. "Computer Vision." Prentice–Hall, Englewood Cliffs, NJ, 1982.
18. MARR, D. "Vision." Freeman, New York, 1980.
19. SIMPSON, P. K. "Artificial Neural System." Pergamon, New York, 1990.
20. WASSERMAN, P. D. "Neural Computing: Theory and Practice." Van Nostrand–Reinhold, New York, 1989.
21. KARSSEMEIJER, N., VAN ERNING, L. J. TH. O., AND EIJKMAN, EG. G. J. Recognition of organs in CT-image sequences: A model guided approach. *Comput. Biomed. Res.* **21,** 434 (1988).

22. RUMELHART, D. E., HINTON, G. E., AND WILLIAMS, R. J. Learning representations by back-propagating errors. *Nature* **323,** 533 (1986).
23. BALLARD, D. H. Form perception using transformation networks: Polyedra. *In* "Advances in Computer Vision" (C. Brown, Ed.), p. 79. Lawrence Erlbaum, Hillsdale, NJ, 1988.
24. SHEPP, L. A., AND LOGAN, B. F. The Fourier reconstruction of head section. *IEEE Trans. Nucl. Sci.* **21,** 21 (1974).