Detection of Suspected Malignant Patterns in Three-Dimensional Magnetic Resonance Breast Images

Essam A. El-Kwae, Joel E. Fishman, Maria J. Bianchi, Pradip M. Pattany, and Mansur R. Kabuka

In this article, a Boolean Neural Network (BNN) is used for the detection of suspected malignant regions in 3D breast magnetic resonance (MR) images. The BNN is characterized by fast learning and classification, guaranteed convergence, and simple, integer weight calculations. The BNN learning algorithm is incremental, which allows the addition and deletion of training patterns without unlearning those already learned. The incremental learning algorithm automatically reduces the training set and trains the network only with those examples estimated to be useful. The architecture is suitable for parallel hardware implementation using available Very Large Scale Integration (VLSI) technology. The BNN was trained by using a set of malignant, benign, and false-positive patterns, extracted by experts, from selected MR studies, by using an incremental learning algorithm. After training, the network was tested by means of a consistency checking test, cross validation techniques, and patterns from actual MR breast images. During the consistency test, the BNN was tested by using the same patterns used for training. The BNN classification accuracy in this case was 99.75%, proving the ability of the BNN to select useful patterns from the training set. Then, a leave one out cross-validation (LOOCV) test was done by using patterns from the training set and the classification accuracy was 90%. Next, an extended training set was created by shifting the original patterns in different directions. A cross-validation test was then performed by dividing the set of patterns into a training and a test set. Classification accuracy was compared to the nearest neighbor classifier. Results showed that the BNN achieved an average of 77% classification accuracy while requiring only 34% of the original training set. On the other hand, the nearest neighbor classifier achieved an accuracy of 57.9% while retaining the whole training set. Another test using actual MR slices different from the training set was done and results compared favorably to a radiologist's findings. Test results show the BNN's capability to detect suspected malignant regions in 3D MR images of the breast. The proposed BNN architecture can save the radiologist a great deal of time browsing MR slices searching for suspected malignancies.

KEY WORDS: breast cancer, magnetic resonance imaging (MRI), neural networks, boolean neural networks

BREAST CANCER is a major cause of death among women. It was expected to be the number one cause of death by cancer in women in 1997.[1] Considerable attention is given to breast cancer because of the potential to reduce the mortality rate by participating in screening programs. Even though mammography is the most common technique for cancer diagnosis, 10% to 30% of women who undergo mammography and who have breast cancer have negative mammograms.[2] Also, mammography can not produce clear images of dense (more glandular) breast tissue, and its effectiveness is limited in women with breast implants.[3] For these reasons, scientists are exploring novel nonionizing imaging technologies including magnetic response imaging (MRI), ultrasound, optical imaging, and other technologies.[3] Of these technologies, MRI and ultrasound have been the most studied to improve the sensitivity of breast cancer detection and staging in certain groups of women. Both methods have shown potential for distinguishing between benign and malignant lesions and in detecting tumors in dense breast tissues.[3]

Magnetic resonance imaging is a medical imaging technique that permits the detailed visualization of soft-tissue structures. MRI has become a standard diagnostic tool in today's clinical settings and its increased availability permits routinely scanning patients to detect a variety of lesions in a noninvasive way. Magnetic resonance of the breast is one of the most promising areas of MRI today. Since the introduction of contrast-enhanced MR imaging, two major achievements of breast MRI over previous MR techniques are improved sensitivity and specificity.[4,5] Sensitivity for breast cancer detection using MRI has been reported to be as high as 98.4%.[6] Bolus injection of gadolinium diethylene triamine penta-acetic acid (Gd-DTPA) is now being used to improve breast cancer detection through identification of the rapid enhancement patterns of malignant lesions. Contrast-enhanced MRI may offer promise as an adjunct test in cases where mammography and physical examination are inconclusive. In addition, some studies have

From the Center for Medical Imaging and Medical Informatics, Department of Radiology, University of Miami, FL.

*Address reprint requests to **Mansur R. Kabuka,** Center for Medical Imaging and Medical Informatics, Department of Radiology, University of Miami, 1150 NW 14th St, Suite 301, Miami, FL 33136.*

0897-1889/98/1102-0001$8.00/0

shown that dynamic bolus contrast MR images are more accurate than mammography or ultrasound in detecting some types of cancer.[7]

However, the accurate interpretation of MRI performed visually by highly trained physicians remains an extremely time-consuming and costly task. Automatic methods for screening MRI studies for signal abnormalities are strongly required. This research focuses on detecting suspected malignant regions in MR images of the breast. It is a step towards a completely automatic approach for accurately locating and estimating the volume of breast tumors from MRI images.

To achieve this goal, a Boolean Neural Network (BNN) was used. The BNN is characterized by fast learning and classification, guaranteed convergence, and simple, integer weight calculations. The BNN learning algorithm is incremental, which allows the addition and deletion of training patterns without unlearning those already learned. The architecture is suitable for parallel hardware implementation by using available Very Large Scale Integration (VLSI) technology. The network was trained by means of a set of malignant, benign, and false-positive patterns, extracted by experts from selected MR studies. An incremental learning algorithm was used for training the BNN. After training, the network was tested by using cross-validation methods and actual MR breast studies.

The rest of this article is organized as follows: an overview of the BNN, its applications, structure, and the concept of Radius of Attraction (ROA) are given in the following section. In the section after that, the architecture of the BNN used and the incremental training algorithm are given. Then, in the next section, the extraction of the training patterns used for this study is explained. Results obtained using the consistency test, the cross-validation tests, and patterns from actual MR studies are given in the "Results" section. In the "Conclusion" we sum up our project and outline future directions for research in the field.

OVERVIEW OF THE BNN

The BNN[8,9] can be used to synthesize any boolean function and also performs efficient classification. The BNN was also used to solve optimization problems.[10,11] A feature recognition BNN has been developed;[9] it recognizes patterns after significant noise, deformation, translation, and scaling. It can be trained to recognize even rotated patterns. The BNN was also used for the design of classifiers[12] such as the Nearest-to-Examplar (NTE) classifier and the Binary K-Nearest Neighbor (BKNN) classifier. Incorporation of production rules into the BNN was introduced and an autonomous approach was proposed for image understanding involving BNN unsupervised segmentation, labeling and a task-specific knowledge-base.[13]

In the medical field, the BNN was used for supervised and unsupervised segmentation of MR brain images.[14] Also, the BNN was used for labeling MR brain images.[15] The BNN proved to be faster than traditional methods and other neural networks. It represents a feasible alternative for existing techniques for medical image labeling. Preliminary results of the use of the BNN to extract suspected malignant regions from MR images of the breast were given in El-Kwae et al.[16]

The BNN consists of layers of neurons connected in a feedforward structure. The neurons are fully connected between layers. Nodes in the BNN use a hard-limiter activation function that is simple to implement in hardware or to simulate in software. The layers in the BNN have simple perceptron connections. During learning, traditional neural networks may never converge or stabilize whereas the BNN converges in a single iteration. Weights are integer-valued and weight calculations are simple. After learning, a new training pattern can be applied without unlearning as in other architectures. The weights of a node directly encode the training pattern. For example, to add a binary training pattern i (of dimension n) to a node k in the network, set the weights to:

$$w_{kj} = 2a_{ij} - 1 \quad \text{for } j = 1, 2, \ldots, n \qquad (1)$$

Where a_{ij} is the value of bit j in the training pattern i and $a_{ij} \in \{0,1\}$, while w_{kj} is the weight associated with bit j of node k in the BNN and $w_{kj} \in \{-1, 1\}$. The output of a node (of threshold Θ_k calculated as described below) is binary (ie, the node does or does not fire) and is calculated as:

$$\text{Output} = \begin{cases} 1 & \text{if} \sum_{j=1}^{n} (w_{kj} b_j) \geq \Theta_k \\ 0 & \text{Otherwise} \end{cases} \qquad (2)$$

Where w_{kj} is the weight associated with bit j of node k in the BNN while b_j is the value of bit j in the input pattern b.

The concept of ROA (Radius of Attraction) is used to classify noisy and distorted patterns. The BNN can be either fixed radius, when the ROA is fixed, or variable radius, when the ROA is varied until a match is found. The BNN training algorithm automatically selects the ROA around each node, as will be explained later. Every examplar node a_k is associated with a programmed threshold Θ_k and a ROA r_k. The programmed threshold is then decremented by the ROA to get the actual threshold.

$$\Theta_k = \sum_j a_{ij} w_{kj} - r_k \tag{3}$$

Where a_{ij} is the value of bit j in the training pattern i to be encoded into node k, w_{kj} is the weight associated with bit j of node k in the BNN and r_k is the ROA of node k.

This leads to all patterns that differ from the stored pattern with a hamming distance of r_k to be classified as belonging to the class of that examplar. The ROA builds hyperspheres of radius r_k with the class programmed threshold at its center. It was proved[14] that the number of patterns (of dimension n) enclosed by a ROA r defined over a certain neuron are:

$$^{n}\Lambda_r = \sum_{i=0}^{r} \binom{n}{i} \tag{4}$$

where n is the dimension of the pattern and r is the ROA.

A generalization of the BNN for handling scalar weighted functions was proposed.[14] In this case, input data is quantized into g gray level scales. These can be represented by m bits where $m = \log_2 g$. The EROA (Euclidean ROA) was introduced because the ROA in this case represents the Euclidean distance. One layer of preprocessing was added to the network whose weights can take fixed values from 2^0 to 2^{m-1}. For the MR breast images, input patterns are subimages extracted from actual MR slices. The features comprising the input patterns are represented by the gray level of the pixels in those subimages. Thus, the words "feature" and "gray level" can be used interchangeably for this application. Because the range of gray levels in the MR images used for this study is from 0 to 255, no additional input data quantization was required. The initial preprocessing layer has 8 input bits with fixed weights ranging from 2^0 to 2^7.

THE BNN ARCHITECTURE AND THE TRAINING ALGORITHM

The BNN Architecture

The architecture of the neural network used is shown in Fig 1. The preprocessing layer described

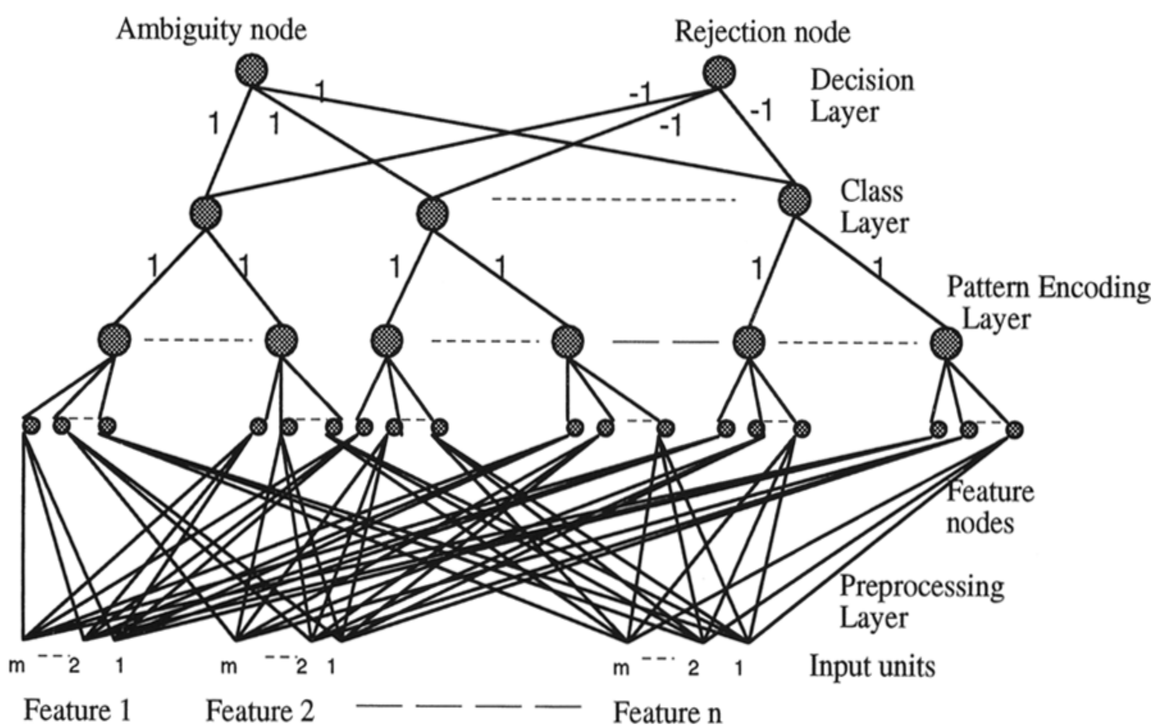

Fig 1. Structure of the BNN used in study.

earlier is shown and is used to input the quantized binarized inputs to the BNN. The network consists of *four* layers: the preprocessing layer, the *pattern encoding* layer, the *class* layer and the *decision* layer. The preprocessing layer is used for quantizing the input values. The number of inputs, at this level, is determined by the desired level of quantization. For example, in MR slices, the range of gray levels is 0 to 255. Thus, 8 bits are required to represent a quantized binarized input value. Weights associated with those inputs range from 2^0 to 2^{m-1} where m is the number of bits used to represent the gray levels of an input pattern. The pattern encoding layer is created incrementally during training while the other two layers are static. A node is added to the pattern encoding layer whenever an input training pattern is misclassified. The weights of the added node directly encode the training pattern.

The ROA of this pattern is calculated as the maximum radius of a hypersphere in which the encoded pattern is at the center such that no two hyperspheres from different classes are allowed to overlap. The output of this node is connected to the single node of the class identification layer representing the training pattern's class. The node fires only if the input pattern is enclosed by the node's hypersphere, ie, the distance between the input pattern and the encoded pattern is less than or equal to the ROA of the node.

The class layer has exactly C nodes, where C is the number of classes in the training set. In this study, C is set to 3 because there are 3 classes in the training set. These classes are *benign, malignant,* and *false-positive.* A node C in the class layer gets its inputs from all nodes representing the class C in the pattern encoding layer. The weights of all connections are set to 1, the ROA to 0, and the threshold θ is set to 1. This node will only fire if one or more nodes belonging to class C in the pattern encoding layer will fire.

The decision layer consists of two nodes, the ambiguity node used to detect if the input pattern was classified to more than one class and the rejection node used to detect if the input pattern was or was not classified by the current network. A pattern is said to be classified by the network if it falls inside one of the hyperspheres of the current nodes in the pattern encoding layer. Both nodes in this layer have exactly C inputs from all the nodes in the class layer. The weights of the ambiguity node are all set to 1, the ROA to 0, and the threshold θ to 2. The weights of the rejection node are all set to -1, the ROA to 0, and the threshold θ to 0. The ambiguity node will only fire if two or more nodes in the class layer fire. The rejection node will fire only if no nodes in the class layer fire.

The BNN Incremental Learning Algorithm

Incremental learning is a multistage learning process used to increase the generalization capabilities of the network. It can also be used to incrementally train a network when the environment changes. The BNN incremental learning algorithm automatically reduces the training set and trains the network only with those examples estimated to be useful. In this case, the teacher need not make an effort to carefully select "good" examples in the training set because the algorithm is capable of automatically selecting those useful examples. Incremental learning also optimizes the number of neurons in the BNN. This is done by setting the ROA around the neuron immediately after an individual exemplar has been learned. If any of the subsequent training exemplars falls within the ROA of one of the existing neurons, it is automatically not learned by the learning algorithm.

Because selected training patterns are directly encoded into the weights of the pattern encoding nodes, learning a new pattern will not affect those patterns already learned. Thus, the BNN can be used in applications where the environment is changing and it is desirable to continuously update the network during operation.

The input to the BNN incremental learning algorithm is a set of N training vectors A_i. The objective is to determine the number of nodes K to represent these input patterns. The output of the algorithm is a set of weight vectors W_K and neuron thresholds Θ_k for each of the K nodes in the network. Initially, a BNN (Fig 1) is constructed. The number of neurons in the pattern encoding layer is set to 0. A counter v_i is associated with each of the available neurons and it is initially set to 0. Training patterns are presented to the BNN. For each pattern, an attempt is made to classify that pattern using the current BNN architecture. If the pattern is correctly classified, there is no need to create a special node to encode the pattern. The counter associated with the firing node is incremented to count the number of patterns in the training set classified to this node. If the pattern is misclassified, the BNN is checked if there are available neurons. In the case where there are excess untrained neurons, the next neuron to be trained is chosen for training with the input pattern. The counter of this neuron is reset after learning takes place. The weight vector of this neuron W_k is calculated as in (1). The neuron's threshold is set as in (3). The ROA is set to half the minimum distance to currently existing neurons from different classes in the BNN, ie, $r_k = 0.5 \min_j D(A_k, W_j)\ \forall j$ such that class label of $W_j \neq$ that of A_i. The ROA of existing neurons from different classes than that of A_i is set to: min(current ROA, 0.5 $D(A_k, W_j)$). This means that adding a pattern might cause the existing hyperspheres to shrink to avoid the overlapping of hyperspheres from different classes. In the case where there are no excess neurons in the network, training of the input pattern has to be done over a preexisting neuron. The least used neuron is selected for training using the counter V_l associated with each neuron. This means that the algorithm discards those nodes that are under utilized or that have poor classification accuracy, allowing the BNN to get rid of noisy patterns. After learning, the V_l of the node used to encode the training pattern is set to 0.

When a pattern is presented to the BNN for classification, whether during training or during actual operation of the BNN after training, the distance between the input pattern and all currently existing hyperspheres is calculated and the pattern is classified to the closest hypersphere. The counter associated with that node is incremented in order to be able to differentiate between active nodes and inactive nodes.

The images used for the BNN analysis were obtained after administration of Gd-DTPA contrast agent. The images were acquired from the whole breast by means of a three dimensional Fourier transform (3DFT), Rf-spoiled gradient echo technique. The signal from the fat in the breast was nulled by using a chemical shift selective Rf pulse applied before each excitation pulse used for image acquisition. The parameters used to acquire images were as follows: TE = 7.0 msec, TR = 43 msec, 1 signal average, 30° flip angle, 3 to 4 mm slice thickness and 22 to 40 slice partitions, 20 to 30 cm field of view, and 256 × 256 image matrix. The total acquisition time was 4 to 5 min with this method.

To test the ability of the BNN to select useful MR breast patterns from the training set and reduce the training set size, an experiment was conducted on a complete MR breast study of 36 images, each of size 256 × 256 pixels. Each pixel assumes a gray value from 0 to 255. Each image was divided into 256 patterns, each pattern being of size 16 × 16 pixels. Thus, the total number of input patterns was 9216. The purpose of the experiment was to find the number of patterns selected by the BNN for different ROAs. The ROA was varied from 0 to 10 and results are summarized in Table 1. Results show that the patterns in MR images form clusters in the feature space. Thus, the ROA approach will be effective in detecting similar patterns and decreasing the training set size.

4. TRAINING SET SELECTION

The images used for this study were T1-weighted, fat suppressed gradient echo images over the whole region of the breast, obtained immediately after the administration of 0.1 mmol/kg Gd-DTPA. The training set was collected from 8 different breast MR studies. Three types of patterns were chosen: benign, malignant, and false-positive. Images having malignant regions were selected, then malignant patterns of size 16 × 16 were extracted. The size of the pattern (16 × 16 pixels) was chosen after studying tumors in many different breast cancer studies. It was found that the whole tumor fits completely in a 16 × 16 pattern or 32 × 32 pattern. The smaller size was chosen to decrease the effect of translation on pattern detection. After selecting malignant regions, an equal number of benign and false-positive patterns was selected from the same image. The number of patterns in the training set was divided as follows: 107 malignant patterns, 106 benign patterns, and 107 false-

Table 1. Effect of the ROA on the Reduction of the Training Set Size

ROA	Selected	Discarded	% Irrelevant
0	6913	2303	24.99%
1	6913	2303	24.99%
2	6840	2376	25.78%
3	6167	3049	33.08%
4	5454	3762	40.82%
5	4944	4272	46.35%
6	4221	4995	54.20%
7	3550	5666	61.48%
8	3051	6165	66.89%
9	2734	6482	70.33%
10	2486	6730	73.03%

positive patterns. A set of 100 training patterns from each class is shown in Fig 2.

During the training session, the network is presented by training patterns. The pattern, originally a matrix, is transformed into a vector by cascading the rows successively. Thus, it becomes a fundamental memory point in the feature space.

To be able to recognize translated patterns, each original pattern of size 16 × 16 was shifted by one pixel in all four directions resulting in an additional set of 8 patterns for each original training pattern. For example, a pattern whose upper left hand corner is located at (x, y) is shifted one pixel in all directions resulting in the following additional patterns: {(x − 1, y − 1), (x − 1, y), (x, y − 1), (x + 1, y), (x, y + 1), (x + 1, y − 1), (x − 1, y + 1)}. Patterns whose corners are located at the edges or corners of their corresponding slices result in a less number of patterns after shifting. Instead of training the BNN using the original patterns, the network is trained using the original patterns in addition to the patterns resulting from shifting. This gives the BNN some immunity against translation, and some degree of deformation and rotation in a window of 2 × 2.[9] The total number of patterns after shifting was found to be 2877 patterns.

The patterns were then used to train the network with the incremental learning algorithm described in section 3. Results of testing the ability of the selected patterns, both the initial and shifted sets, to extract suspected malignant regions in 3D MR images of the breast, are given in the following section.

TESTING

Consistency Test

To test the ability of the BNN to remember the training patterns, a consistency test was conducted. The network was first trained using 2000 patterns selected randomly from those in the shifted training set (2877 patterns). The BNN was then tested with the same patterns used for training. The output classifications from the network were compared to the desired results. The BNN correctly classified 1995 patterns out of 2000 (99.75%). These results confirm the ability of the network to remember all training patterns.

Test by Cross-Validation Methods

To test the generalization ability of the neural network to patterns that were not included in the training set, cross-validation methods were employed. In these methods, some of the patterns are selected at random from the training set for training the network, and the remaining patterns are used for testing the trained network. The output values are compared with the expected output. Various combinations of training and testing pairs can be selected by using a random number generator. In some of the cross-validation tests conducted in this study, the leave-one-out cross-validation (LOOCV) test was used. In this case, the whole set of training patterns is used for training the network except a single pattern which is then submitted to the network for classification. This process is repeated in a round-robin fashion to all patterns in the training set.

In the first LOOCV test, the network was trained using the initial set of 320 patterns extracted from MR studies. Classification results are summarized in the Table 2. The intersection of each row and column gives the number of patterns from the class designated by the row label that were classified as belonging to the class designated by the column label. For example, there were 102 benign patterns

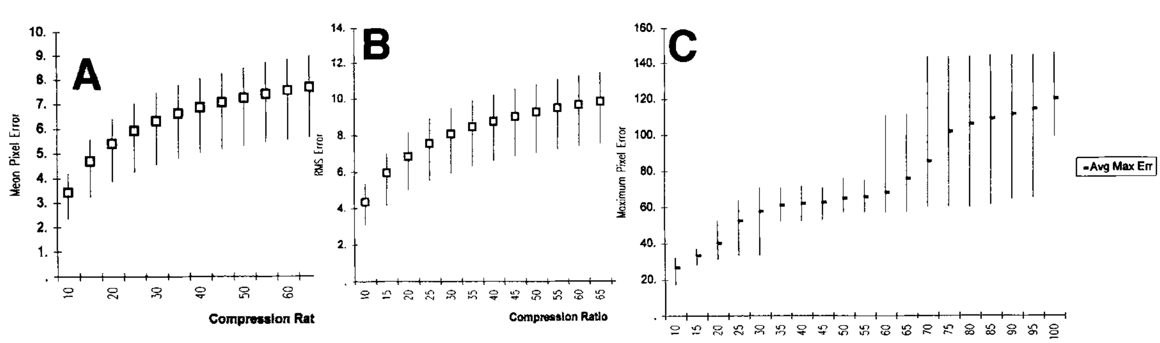

Fig 2. Samples from the BNN training patterns. A set of 100 training patterns (16 × 16 pixels) is shown for each of the three predefined classes: (A) False-positive, (B) Benign, (C) Malignant.

Table 2. Classification Results for the Initial Training Set When Applying LOOCV

	Benign	Malignant	False-Positive	Total
Benign	102	0	4	106
Malignant	13	80	14	107
False-Positive	58	6	43	107

that were classified as benign by the BNN whereas 4 originally benign patterns were recognized as false-positive.

Table 3 shows the classification accuracy for this test. Of the 287 correctly classified patterns, 62 (21.6%) were classified as benign instead of false-positive and vice versa. This was not considered to be a wrong classification because both types are negative patterns. Classification accuracy was about 90%, attesting to the quality of the extracted patterns.

In the context of this work, sensitivity is defined as the number of malignant patterns detected by the BNN divided by the number of malignant patterns in the test set. Specificity is defined as the number of negative patterns detected by the BNN divided by the true number of negative patterns in the test set. Based on these definitions, the sensitivity of the BNN for this test was about 75% whereas the specificity was about 97%.

In the second cross-validation test, the shifted training set was divided into two sets: a training set of 2000 patterns selected randomly, and a testing set of the remaining 877 patterns. The BNN was created as described in section 3 and trained using the incremental learning algorithm described in section 3. The test set was then submitted to the BNN and classification accuracy was then calculated. Results were then compared to the nearest neighbor classifier. The nearest neighbor (NN) classifier assigns to an unclassified pattern the classification of the nearest of a set of previously classified patterns. The NN is non-parametric, or distribution free, ie, it does not depend on any assumptions about the underlying joint distribution of the input samples. An excellent resource for the NN is Dasarathy's book.[17] Cover and Hart[18] proved that the probability of error R of the NN rule for a

Table 3. Classification Accuracy for the Initial Training Set When Applying LOOCV

Number of Patterns	320
Correct (Correct %)	287 (89.69%)
Wrong (Wrong %)	33 (10.31%)

Table 4. Training Results of the BNN on the Shifted Training Set and Comparison of its Classification Accuracy to the NN Classifier

Test Number	BNN Size	BNN Correct	NN Correct
1	690	685	519
2	701	666	505
3	725	701	523
4	698	680	535
5	698	675	501
6	683	658	495
7	701	681	495
8	707	682	498
9	682	690	505
10	691	673	499
11	702	661	505
12	700	661	515
13	677	654	500
14	707	693	520
15	712	699	533
16	685	664	478
Average	697.44	676.44	507.86
SD	12.45	14.85	15.17
Average %	34.87%	77.13%	57.91%

large number of samples M is:

$$R^* \leq R \leq R^*\left(2 - \frac{MR^*}{M-1}\right) \qquad (5)$$

Where R* is the Bayes probability of error (the minimum error probability of all decision rules and can be achieved when the underlying probability distribution is known and taken into account). Thus, for any number of classes, the nearest neighbor rule is bounded by twice the Bayes probability of error. To achieve the bounds, it is necessary to store a large number of samples of known classification.[18,19] In general, NN algorithms play an important role in inductive machine learning because of their simplicity and their ability to give highly accurate predictions after a short learning phase.[20] NN queries have also been utilized in diverse applications including medical images.[21]

The second cross-validation test was repeated 16 times. Results of each of the 16 trials were reported in Table 4. This was necessary because the cross validation test selects the training and testing patterns randomly from the whole set of patterns.

Table 5. Classification Results for the Shifted and Reduced Training Set when Applying LOOCV

	Benign	Malignant	False-Positive	Total
Benign	9	0	16	25
Malignant	8	336	15	359
False-Positive	62	4	227	293

Table 6. Classification Accuracy for the Shifted and Reduced Training Set when Applying LOOCV

Number of Patterns	677
Correct (Correct %)	650 (96.01%)
Wrong (Wrong %)	27 (3.99%)

Owing to this randomness, each run of the cross-validation test might give different classification accuracy. To decrease the statistical variation in the results, the test was repeated 16 times. Results for each test were reported in addition to the average and standard deviation. This gives a more reliable assessment of the accuracy of classification. The column labeled "Test Number" refers to the order of the test. The column labeled "BNN size" refers to the number of nodes required for the pattern encoding layer in each case. The columns labeled "BNN Correct" and "NN Correct" correspond to the correct classification of the 877 test patterns by the BNN and the nearest neighbor (NN) classifier, respectively. Results show that the BNN incremental training algorithm was able to reduce the original training set to about 35% of its original size and, in the mean time, improved generalization ability about 20% than the NN classifier.

In the third cross-validation test, the reduced training set resulting from applying the BNN incremental training algorithm to the training set of 2000 patterns was to be tested. A LOOCV test was conducted on the 677 patterns selected by the BNN in round 13 of the previous test. This particular set was used because it gave the least BNN size, and hence least cost of implementation, of all 16 tests conducted. Of the 677 patterns selected by the BNN, 25 patterns were benign, 359 were malignant, and 293 were false-positive patterns. Classification results are summarized in Table 5. The intersection of each row and column gives the number of patterns from the class designated by the row label that were classified as belonging to the class designated by the column label. For example, there were 336 malignant patterns that were classified as malignant by the BNN while 15 originally malignant patterns were recognized as false-positive.

Table 6 shows the classification accuracy for this test. Of the 650 correctly classified patterns, 78 (12%) were classified as benign instead of false-positive and vice versa. Again, this was not considered to be a wrong classification. The percent of correctly classified patterns increased from about 90% in the first LOOCV test on the original training set to about 96% on the shifted, then reduced training set. This implies the robustness of the new generated training set. Based on the earlier definitions of sensitivity and specificity, the sensitivity of the BNN for this test improved to about 94% whereas the specificity also improved to about 98%.

Testing on MR Studies

The cross-validation tests showed that the network has generalization, shift invariance, and training set reduction capabilities. This suggested test-

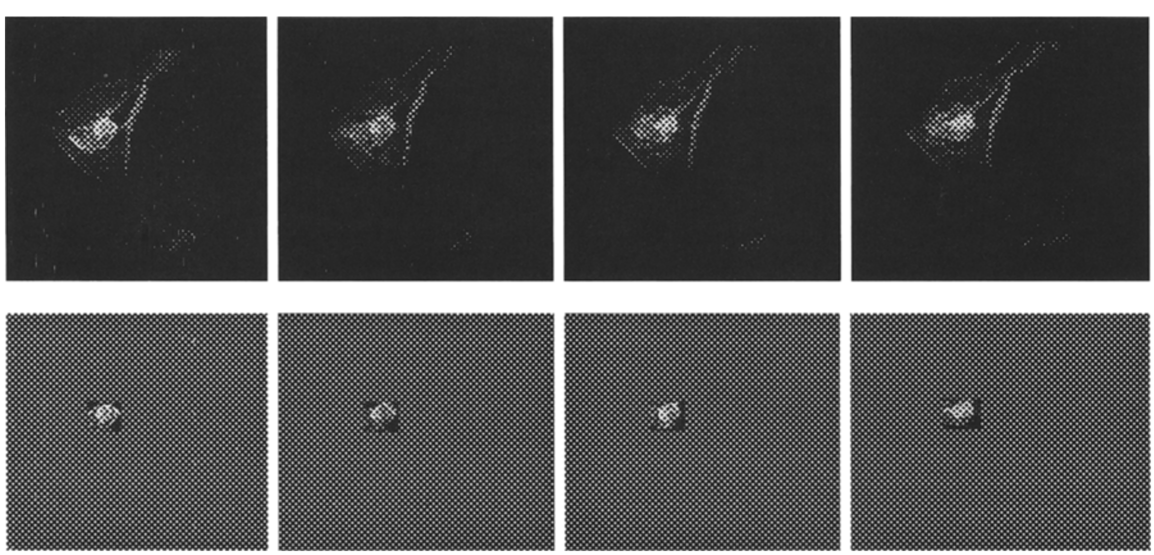

Fig 3. Testing with a series of MR slices. The top set of images are the original images, the bottom images are the images after extracting the suspected malignant regions only.

ing the network on actual MR studies to detect suspected malignant regions. Each image in the test study was divided into 16 × 16 patterns and fed to the network for classification. The patterns detected to be malignant were kept whereas all other benign and false-positive patterns were blackened, ie, the goal was to create new images showing only the suspected malignant patterns in the image. Eight cases were used for this test. In each case, the BNN network was trained by using patterns extracted from the other 7 studies. This shows the ability of the BNN to generalize to patterns that were not seen during training. The BNN findings were approved by experts. Results of extracting malignant regions from a series of MR slices in three of the test studies are shown in Figs 3 to 5.

In general, the network was able to extract suspected malignant regions efficiently. In some cases, the BNN classified an isolated pattern as being malignant. Such patterns can be easily discarded by using a postprocessing step. Any classified malignant patterns that do not have any other malignant patterns in their neighborhood, either the 2D or the 3D neighborhood, may be discarded. More rules can be added and a simple rule base can be used to enhance the robustness of the suspected malignancies extracted by the BNN.

Because the BNN learning algorithm is incremental, more patterns and even more classes of patterns can be added to the network after training. For example, malignant patterns can be divided into different sub-classes according to the size or any other diagnos-

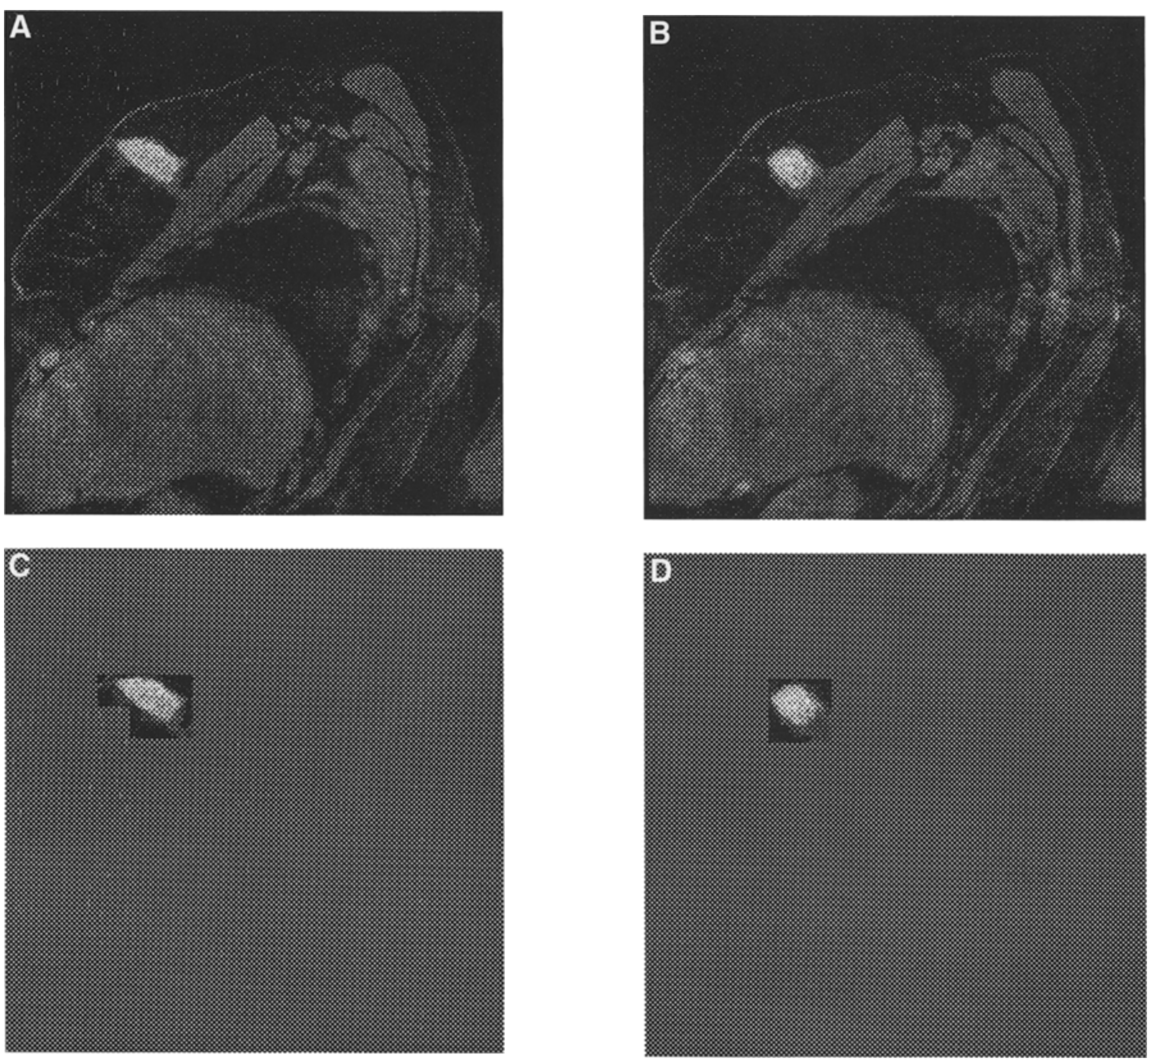

Fig 4. Testing another set of MR slices. (A, B) The original images. (C, D) The images after extracting the suspected malignant regions.

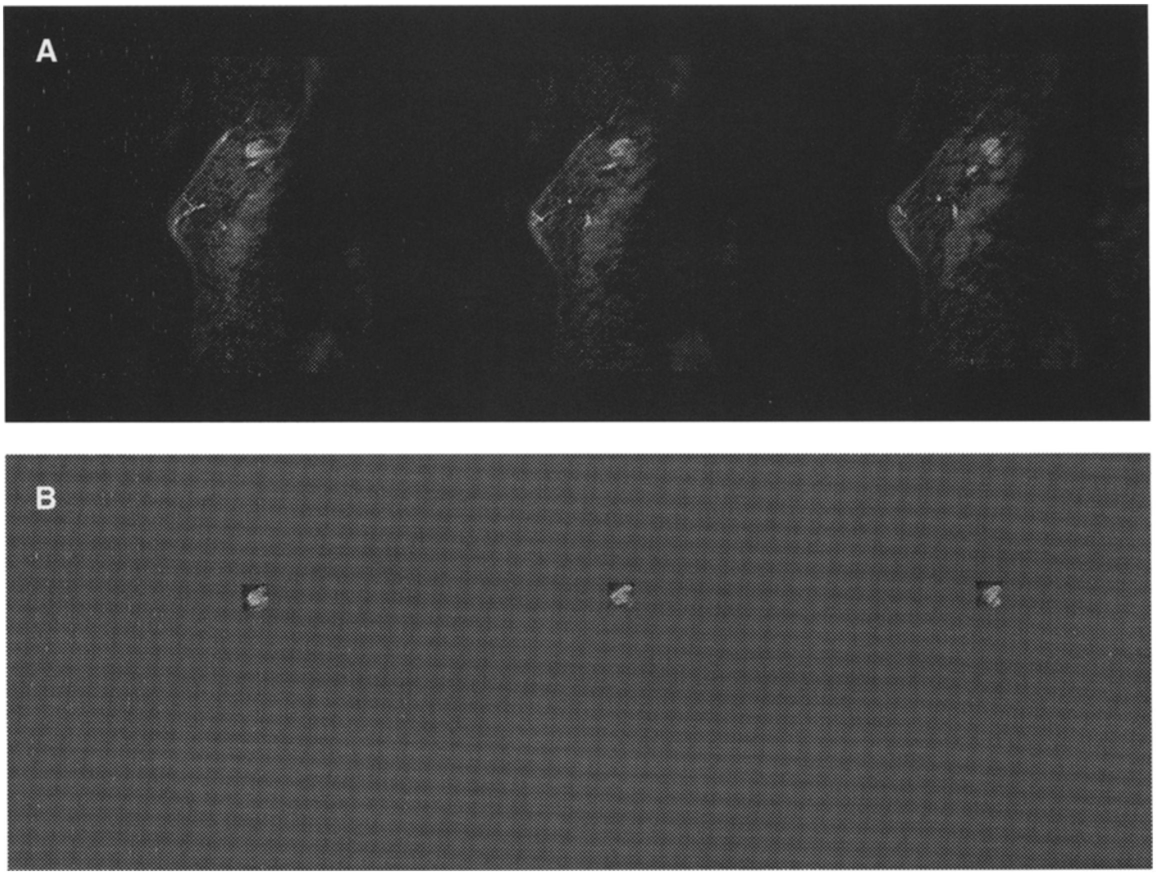

Fig 5. Testing a third series of MR slices. (A) The original images. (B) The images after extracting the suspected malignant regions.

tic factor. Also, a new class of glandular patterns can be added to be able to differentiate between glands and malignancies in dense breast tissues.

6. CONCLUSION

A Boolean Neural Network (BNN) was used for the detection of suspected malignant regions in 3D breast MR images. The network was trained by using a set of malignant, benign, and false-positive patterns, extracted by experts from selected MR studies, using an incremental learning algorithm. The algorithm was enhanced by training the BNN on shifted patterns in addition to the extracted patterns. This gives the BNN some immunity to shifting and deformation. Cross-validation testing techniques showed the robustness of the selected training set and the ability of the BNN learning algorithm to select a compact training set. Comparison with the nearest neighbor classifier showed the improved generalization capabilities of the BNN learning algorithm. The reduced training set selected by the BNN from the shifted training set greatly improved the robustness of the training set on the expense of increased BNN size. The network was tested using complete images from 3D MR breast studies. Results show that the BNN has good generalization, shift invariant capabilities, and a great aptitude to extract suspected malignant regions. Simple postprocessing rules can be used to improve on the results obtained by the BNN. The BNN incremental learning algorithm allows the addition of new patterns, and even new classes, at any point of time without any need for retraining.

Further research is needed to determine the adequate size of the pattern database to be used for this network. One way of doing that is to build the network incrementally. The BNN is tested by using any newly available MR breast images. If it fails, the network is switched to the training mode, and then taught by the malignant patterns it failed to recognize. This can be done until the BNN achieves a certain required level of accuracy.

REFERENCES

1. Guerin S, Laplanche A: Statistics of mortality in 1994 and predictions of death caused by cancer 1997. Presse Med 26:1149-1153, 1997

2. Giger ML: Computer Aided Diagnosis, in A Categorical Course in Physics, Technical Aspects of Breast Imaging. Oak Brook, IL, RSNA, 1993

3. *National Cancer Institute:* CancerNet: Cancer Facts. Bethesda, MD, National Institutes of Health, 1995

4. Heywang-Koebrunner S, Viehweg P: Sensitivity of contrast-enhanced MR imaging of the breast. Magn Reson Imaging Clin N Am 2:527-538, 1994

5. Piccoli C: The specificity of contrast-enhanced breast MR imaging. Magn Reson Imaging Clin N Am 2:557-571, 1994

6. Weinreb J, Newstead G: Controversies in breast MRI. Magn Reson Q 10:67-83, 1994

7. Gilles R, Guinebretiere JM, Shapeero LG, et al: Assessment of breast cancer recurrence with contrast-enhanced subtraction MR imaging: Preliminary results in 26 patients. Radiology 188:473-478, 1993

8. Hussain B, Kabuka MR: Neural net transformation of arbitrary boolean functions. Presented at SPIE 1992 International Symposium on Optical and Applied Science: Neural and Stochastic methods in Image and Signal Processing, San Diego, CA, December 1992

9. Hussain B, Kabuka MR: A novel feature recognition neural network and its application to character recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence 16 SPIE 1766:355-367, 1993

10. Shirish B, Nigel J, Kabuka MR: A boolean neural network approach for the traveling salesman problem. IEEE Transactions on Computers 42:1271-1278, 1993

11. Shirish B, Nigel J, Kabuka MR: A real-time solution for the traveling salesman problem using a boolean neural network. Proceedings of the 1993 IEEE International Conference on Neural Networks, San Francisco, California, 1993, pp 1096-1103

12. Srinivas G, Kabuka MR: Design of supervised classifiers using boolean neural networks. IEEE transactions on Pattern Analysis and Machine Intelligence 17:1239-1246, 1995

13. El-Kwae EA, Kabuka MR: A boolean neural network approach for image understanding. Proceedings of Artificial Neural Network in Engineering Conference (ANNIE '96), St Louis, MO, Nov 10-13, 1996, pp 437-442

14. Shirish B: Image segmentation with boolean neural networks. Masters thesis, University of Miami, 1994

15. Xiahong L, Shirish B, Kabuka MR: Labeling of MR brian images using a boolean neural network. Proceedings of Symposium of Computer Assisted Radiology, Winston-Salem, NC, June 1994

16. El-Kwae EA, Bianchi MJ, Kabuka MR: Detection of malignant regions in 3D MR breast images using a boolean neural network. Symposium for Computer Applications in Radiology (SCAR 96), Denver, CO, June 6-9, 1996, pp 304-309

17. Dasarathy BV: Nearest Neighbor (NN) Norms: NN Pattern Classification Techniques. New York, NY, IEEE Computer Society Press, 1991

18. Cover TM, Hart PE: Nearest neighbor pattern classification. IEEE Transactions on Information Theory 13:21-27, 1967

19. Tou JT, Gonzalez RC: Pattern Recognition Principles. Reading, MA, Addison-Wesley, 1974

20. Wettschereck D, Dietterich TG: An experimental comparison of the nearest-neighbor and nearest-hyperrectangle algorithms. Machine Learning 19:5, 1995

21. Korn F, Sidiropoulos N, Faloutsos C, et al: Fast Nearest Neighbor Search in Medical Image Databases. Presented at Conference on Very Large Databases (VLDB), Bombay, India, Sept 1996