Present and future methods of mine detection using scattering parameters and an artificial neural network

Gregory L. Plett, Takeshi Doi,

Durand 104, Department of Electrical Engineering,
Stanford University, Stanford, CA 94305-4055

Don Torrieri,

Army Research Laboratories,
2800 Powder Mill Road, Adelphi, MD 20783

ABSTRACT

The detection and disposal of anti-personnel landmines is one of the most difficult and intractable problems faced in ground conflict. This paper first presents current detection methods which use a separated aperture microwave sensor and an artificial neural-network pattern classifier. Several data-specific pre-processing methods are developed to enhance neural-network learning. In addition, a generalized Karhunen-Loéve transform and the eigenspace separation transform are used to perform data reduction and reduce network complexity. Highly favorable results have been obtained using the above methods in conjunction with a feedforward neural network. Secondly, a very promising idea relating to future research is proposed that uses acoustic modulation of the microwave signal to provide an additional independent feature to the input of the neural network. The expectation is that near-perfect mine detection will be possible with this proposed system.

Keywords: landmine detection, neural network, pattern classification, feature selection, microwave sensor, acoustic sensor

1. INTRODUCTION

Anti-personnel landmine detection is a problem of significant military, economic and humanitarian concern. Not only do landmines pose a deterrent to military activity during conflict, they also remain lethal long after that conflict is over. The casualty statistics are horrifying; eighty percent of the victims of anti-personnel mines are civilians, many of whom are children.[1] The number of such mines currently in place is estimated at over 100 million, and is increasing at the rate of one half to one million mines per year. There is great motivation to research new countermine systems. Current anti-personnel land mines cost between \$3–\$25 per mine and clearance methods cost between \$300–\$1000 per mine removed.[2] A mine detection system which will operate with a high degree of accuracy while being simple enough to deploy at low cost is needed.

One of the problems faced by the countermine system is that newer mines contain very little metal, and thus it is not feasible to use conventional metal detectors to locate them. A more sophisticated approach is presented here, where neural networks are trained to detect simulated land mines using data collected from a mine lane facility located at Ft. Belvoir, VA. This data consists of a set of measurements made with a Separated Aperture Sensor operating in the "waveguide near cutoff" mode.[3] Previous work has shown that neural-network pattern classifiers can successfully detect anti-tank mines,[3–5] and the present work addresses the much more difficult problem of detecting smaller anti-personnel mines.

Emphasis is placed on the processing of the sensor data for use with a neural-network classifer in order to maximize detection rates while minimizing false-alarm rates. The structure and type of data used for the input to the network is explained and justified. Three methods, which mitigate certain undesirable and uncontrollable properties of the training data set, are discussed. These are: data flattening, outlier removal, and data balancing. These procedures permit better neural-network learning and improved generalization. In addition, a generalized Karhunen-Loéve transform and the eigenvalue separation transform are explored as effective preprocessing methods to perform data reduction and increase mine detectability. The performance of artificial neural-network mine detectors incorporating these methods are summarized. Finally, a novel feature to be introduced to the data gathering aparatus is discussed.

2. SENSOR DESCRIPTION AND DATA ACQUISITION

An experimental apparatus at Ft. Belvoir was used to obtain the training and test data from a simulated mine lane. The fundamental operating principles of the separated aperture sensor are described in references 3 and 6. A block diagram of the measurement system is shown in Fig. 1. A Hewlett Packard 8753A Network Analyzer capable of generating sinusoidal signals between 300 kHz and 3 GHz was attached, through an HP 85046B S-parameter test set, to a separated aperture sensor. The sensor, located (nominally) 2" above the ground, consisted of a transmitting aperture or horn which injected microwaves into the soil, and a receiving aperture which provided a return signal to the Hewlett Packard equipment. For the data set used in the current work, the resonant frequency of the horns was 1 GHz. Figure 2 shows a pictorial diagram of the sensor. The horns look like metal troughs with the dipole antennae running longitudinally.

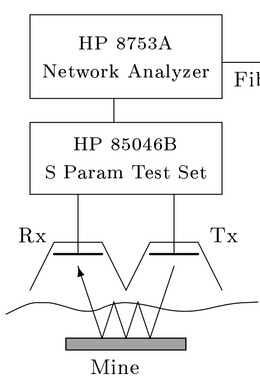

Figure 1: Block diagram of the measurement setup.

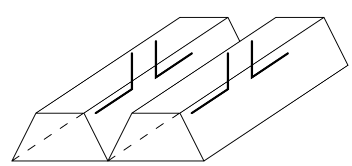

Figure 2: Diagram of the microwave horns and antennae.

Both the signal received by the "receiving" horn and the return signal picked up by the transmitting aperture are sensed. The HP S-parameter test set automatically records the complex-valued scattering parameters or S-parameters of the two-port network formed by the microwave horns and the earth waveguide. The tests use continuous-wave sine waves, and standing waves were detected and used to compute the S-parameters. The horns were connected directly to the S-parameter test set, and the measurements were repeated five times and averaged to reduce noise.

The energy radiated from a dipole antenna is transmitted in both electric and magnetic waves. It has been found that the leakage mode which propagates across the ground surface is a TE mode wave[7]; the detection mode, however, is largely transmitted through the soil over the target. There is some energy absorbed by the target, with the strongest absorbed field component parallel to the dipoles.

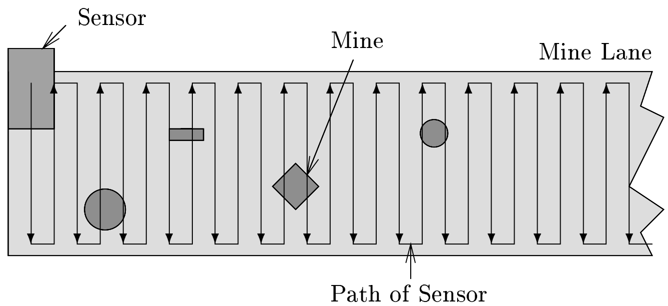

Figure 3: The movement of the sensor over the mine lane during data acquisition.

The mine lane facility at Ft. Belvoir was 790.5" long and 40.5" wide. Data was collected every 1.5" across the length and width of the lane, thus forming a 527 row by 27 column grid of measurements. At each point, all S-parameters were measured (S_{11},S_{12},S_{21},S_{22}). Eleven different frequencies in 40MHz intervals were used ranging from 800MHz to 1200MHz, inclusive. At each point, the magnitude and phase component of each S-parameter was measured for a total of 88 ($4\times11\times2$) measurements at each location in the grid. Figure 3 shows how the sensor head was scanned over the mine lane.

Previous data sets were acquired using a wheeled cart which housed the data collection equipment. The more recent data was collected by a system hung on an overhead rail. The measuring system was controlled at a distance by an HP 9000 computer via optical fiber (HP-IB) link. The computer was able to move the sensor head laterally with a worm gear control and record all measurements automatically. To move the sensor longitudinally (from row to row), a human operator moved the apparatus the required 1.5" along the overhead rail. In addition, the new data set was measured over a mine lane of sandy soil which included five different mine types:

1. 6" diameter × 3" wood disk.
2. 5" diameter × 3" hard plastic disk.
3. 3.5" diameter × 3" nylon disk.
4. 6" × 6" × 3" wood block.
5. 4" × 2" "butterfly" shaped plastic and metal target.

There were a total of 114 targets: 15, 30, 12, 15 and 42 of each mine type, respectively.

3. NEURAL NETWORKS FOR PATTERN RECOGNITION

Computation by means of networks of artificial "neural" elements is an idea that has proven itself capable of solving problems that would be impossible or impractical to solve by other means. For an excellent tutorial paper on neural networks, the reader is referred to reference 8. In short, the following are valuable attributes of neural networks which are not generally shared by other computing systems:

- Neural networks are not *programmed.* Rather, they are *trained* to respond and perform highly refined tasks for which no precise rules exist. One such task is pattern recognition (such as the discrimination of mines in a minefield), in situations where no empirical or theoretical rules are known.
- Neural networks can generalize, that is, deliver accurate responses to inputs that were not presented during training. This is a crucially important feature since no system can be trained on all possible mine and background patterns.
- Neural networks can give very high-speed response to input stimuli, once trained, due to their inherently parallel nature.
- Neural networks can be continually trained, keeping up with changing input environments.

Neural networks are interconnected structures of simple processing elements which crudely model the function of a biologic neuron. Each artificial neuron (hereafter referred to simply as neuron) has the composition shown in Fig. 4. Internally, the scalar product of the input vector* and a weight vector is computed, and the output is a non-linear function of this scalar product: $out = F(X \cdot W)$. In our work, $F(\cdot) = \tanh(\cdot)$. By modifying the value of the weight vector, different output functions can be realized.

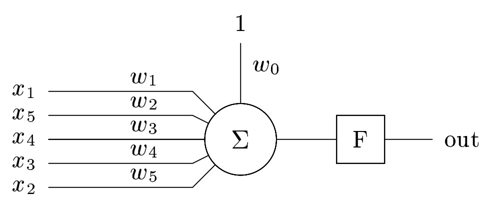

Figure 4: Structure of a neuron.

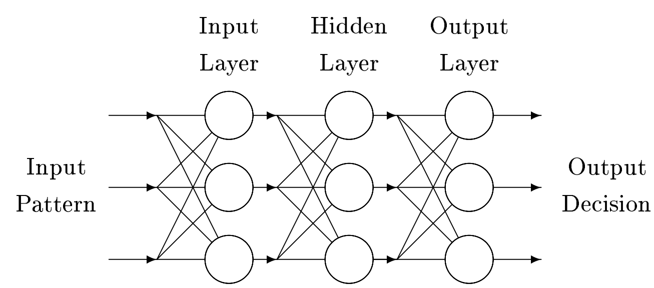

Figure 5: Neural-network structure.

By combining many of these simple neurons in a layered network, where one element of the input vector of each neuron in a layer is connected to the output of all neurons on the previous layer, a very powerful computational tool is achieved. This structure is show in Fig. 5. It has been shown by Kolmogorov that such a network with a single hidden layer, and a sufficient number of neurons is capable of computing (with some set of weight vectors) any continuous nonlinear function. Thus, for example, it is able to compute a discriminant function used for pattern recognition.

To be useful, we need to be able to compute the weight vectors which can be used to create the pattern discrimination function. Several algorithms exist for this purpose, the most popular of which is the "Backpropagation" algorithm. This algorithm requires a set of data which includes the input patterns and known classes for each of these patterns. By iteratively presenting patterns and desired responses to the network, and updating the weights using the backpropagation algorithm, the network learns to correctly classify the patterns in the training set. The network is also able to generalize, and correctly classify many patterns which are not in the training set.

4. MINIMAL WINDOW

In the initial research, each training and test pattern applied to the neural network was obtained by using a square spatial window of data from many nearby points. These spatial windows ranged in size from 5×5=25 points to 13×13=169 points. At each point in a window, only the S_{11} parameter for each of a number of frequencies

*The input vector is augmented by adding a zeroth element, always equal to 1 was included in the pattern. The decision as to whether or not a mine was actually present at the center of the window was based on data collected over the entire window area.

Although the use of spatial windows yielded good results, their application to a practical mine detector presents great difficulties because of their inherent operational requirements. Consequently, the spatial windows were abandoned. Instead, each pattern applied to the neural network was obtained from data measured at a single point. This new approach could be interpreted as the use of a minimal (1×1) window. The principal advantages of a minimal window are:

1. Easier and faster measurements,
2. The possibility of mine detection near natural or man-made obstacles encountered in the field,
3. The possibility of collecting larger sets of training and test patterns by including points near the edges of the mine lanes,
4. A reduction in the complexity of the neural network,
5. Better accommodation of irregular terrain.

A minimal window requires that each input pattern applied to the neural network be derived from measurements taken from a single location. These measurements potentially include the magnitudes and phases of all four scattering parameters at a number of frequencies. However, extensive simulation trials have indicated that little or no performance improvement is obtained by including the phase measurements or by including measurements at more than seven frequencies. Thus, each input pattern can be reduced to only 28 components—the magnitudes of the four scattering parameters at the seven central frequencies.

5. DATA FLATTENING

One significant deterrent to automated mine detection was found to be caused by large, step-like jumps in magnitude which periodically occur between successive rows of a data set. Figure 6 shows one example jump in the 800 MHz S_{11} data between rows 414 and 415.

It is thought that these jumps are caused by a change in the ambient temperature of the mine lane from one day to the next. At the beginning of the day, the temperature would be significantly different from the evening before, and a jump would occur between rows in the data. In addition to jumps, there is a significant downward slope in the data, which means that even a short lunch break could be the cause of a jump in the data (as evidenced by some jumps "down"). Unfortunately, not all of these jumps are detectable or removable.

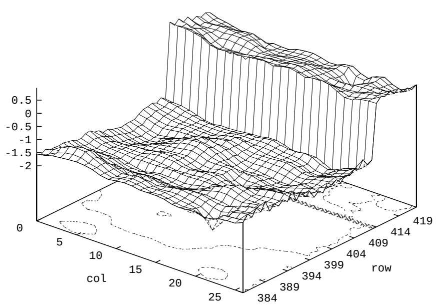

Figure 6: Unflattened data.

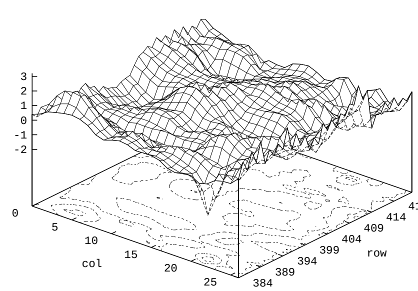

Figure 7: Flattened data.

The jumps are artificial in the sense that they would not occur in a field system searching for mines. Furthermore, it is hypothesized that the jumps decrease the network's performance since many different signal levels artificially correspond to mine patterns. Therefore, it is desirable to remove the jumps from the data if possible.

A preprocessing function attempts to eliminate the jumps. First, the mean and standard deviation of both the magnitude and phase components of each data set† are computed. Next, the average magnitude and phase value for each row of data are computed. If the difference in the average of either the magnitude or phase between two successive rows is greater than one standard deviation of the data set, a jump is hypothesized to have occurred at this location. One further restriction is used to distinguish between jumps and true high-variance changes in the data. Once a jump is postulated to have occurred, it is verified by checking that all the columns change value in the same direction across the jump.

The ramps, and thus the jumps at the ramp edges, are removed as follows: for each interval of data between two jumps a linear regression calculation of the best least squares fit of a line to the column-averaged data is made. This is the approximation of the "ramp" experienced during measurement. This line is extended over all of the columns to form a plane, and is subtracted out of the data.

For example, the 800 MHz S_{11} data file has an original average row magnitude as displayed in Fig. 8. Many jumps are clearly evident; consider the one between rows 414 and 415 (this is the same jump as in Fig. 6). When the jumps are removed using this method, the new average row magnitude is as shown in Fig. 9. Looking at the data in three-dimensional format (Fig. 7) we see that the jump has been removed.

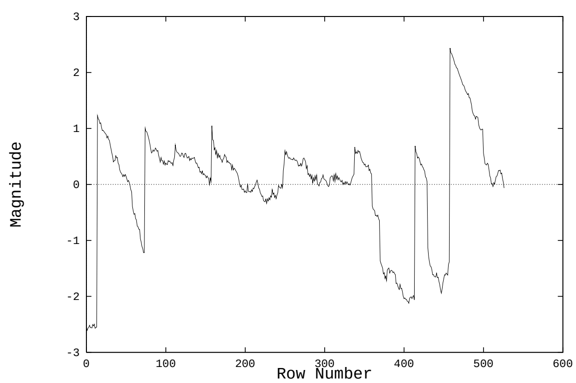

Figure 8: Unflattened average row magnitudes.

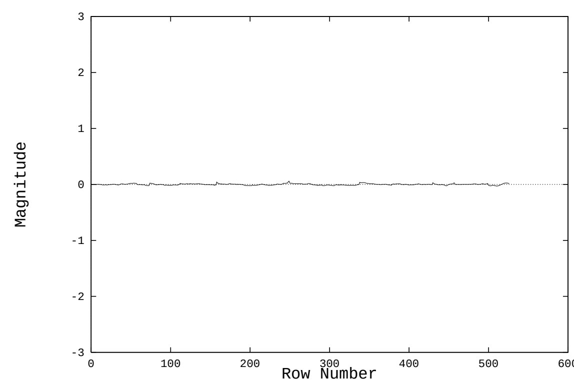

Figure 9: Flattened average row magnitudes.

6. OUTLIER REMOVAL AND BALANCED DATA SETS

The supervised learning of a neural network requires the association of a desired response with each training and test pattern. An assignment of a desired response was possible for each tested location in the mine lane because the size, orientation, and center positions of the mines were known. At each location used to measure a training pattern, a positive response was assigned if a mine overlapped the 1.5"×1.5" grid square with that location at its center. A negative desired response was assigned if the grid square was not overlapped by a mine. A training or test pattern with a positive desired response is called a mine pattern, and one with a negative desired response is called a background pattern.

The generalization capability of neural networks can be impaired by overfitting noisy training patterns. The separating surface generated by a neural-network classifier may exhibit excessive fluctuations rather than the smooth boundary needed for good generalization. To remove the statistical outliers which cause the overfitting

†The term "data set" here refers to 527×27 magnitude and phase measurements made for one S parameter and frequency combination.

from the training set, two approaches may be considered. In the first approach, vector quantization or some other mathematical clustering technique is used to find anomalous patterns and then either remove them from the training set or suppress their influence on the separating surface. The second approach entails an analysis of the physical conditions that may cause certain patterns to be outliers. Both methods were pursued in this work, with the finding that the much simpler second method consistently produced better results. Details of this latter approach are presented here.

Outliers are most likely to be generated when the sensor is placed near the edge of a mine, primarily because of the potential multipath effects. For example, at a mine boundary, a negative desired response may be assigned based on the sensor position, even though the training pattern is much more typical of a mine pattern or is much different from the other training patterns. If all the training patterns measured near mine edges are regarded as potential outliers and are removed from the training sets, one would expect improved neural-network generalization.

To eliminate potential outliers, the pattern measured at a point was retained in the training set only when the four nearest surrounding points had the same desired response. Along the edges and corners of the mine lane, a missing neighbor was excluded from the requirement. This method of removing outliers is called the four-nearest-neighbor outlier removal method. More nearest neighbors can be included in the criterion, which leads, for example, to the eight-nearest-neighbor method.

An immediate problem resulting from the use of the nearest-neighbor methods was that the number of mine patterns was significantly reduced, thereby impairing the generalization capability of the neural network. To prevent this problem, only potential background patterns were subjected to the nearest-neighbor methods. All mine patterns were retained in the training set.

The scarcity of training patterns measured in the vicinity of mine edges may cause a mine detector to be insensitive to the presence of a mine when the sensor is placed above the edge of a mine. However, in a practical implementation, the mine detector will probe the central part of the mine at some time during the detection process. Therefore, the exclusion of statistical outliers will not degrade a practical implementation.

The backpropagation algorithm used for neural-network training converges very slowly, if at all, for binary classification problems in which the bulk of the training patterns belong to a single class. In the presence of such a data imbalance, one approach to accelerating the convergence is to modify the backpropagation algorithm.[9] A much simpler approach, which was adopted, is to artificially create a balanced training set of approximately equal numbers of mine patterns and background patterns. The balanced training set is created by replicating each pattern in the smaller class, the class of mine patterns, a sufficient number of times. As a result, the mine and background patterns are presented to the neural network with approximately the same frequency during the training phase. The use of a balanced training set together with the nearest-neighbor outlier removal method produced a substantial improvement in the simulated performance of the mine detector relative to its performance using the original imbalanced set.

7. THE GENERALIZED KARHUNEN-LOEVE TRANSFORM

The Karhunen-Loéve (KL) transform[10] is a linear transformation of the form:

$$Y = A^T X, \tag{1}$$

where A is the orthonormal transformation matrix, and X is a k-dimensional input pattern vector. The motivation behind using this transformation on the mine lane data is twofold. First, the pattern space may be transformed in such a way that a network is more capable of separating the mine patterns from the background patterns. Secondly, data reduction can be achieved, that is, we can compute $Y_m = A_m^T X$ where A_m contains the first $m \leq k$ columns of A, and Y_m has the resulting dimension m. Generally, better separation and lower dimensionality can be attained at the same time, and the data reduction is achieved without a significant loss of "energy" from the input data, and thus without a significant loss of performance. This data reduction will result in faster training and faster and more compact network realizations.

For the KL transform, the matrix A is set to be the modal matrix of X. This matrix contains, as its columns, the eigenvectors of the covariance matrix of pattern vectors. The columns are ordered by descending eigenvalue, and in order to maximize the energy retained by the transformation, we retain those columns corresponding to the largest eigenvalues.

For pattern recognition, though, the metric we should use is not that of maximizing the preserved energy of the transformation, but rather maximizing the average distance between patterns of the two classes. Toward this end, we have generalized the concept of the KL transform such that A contains the eigenvectors of the modal matrix of background patterns only. This transform then produces components of each pattern of the two classes in such a way that the components of one class have a larger average energy in certain directions than the components of the other class. In principle, the correlation matrix calculated from only the mine patterns could be used instead, but this approach was not adopted because of the relative scarcity of mine patterns and the consequent potential inaccuracy in the computations.

8. THE EIGENSPACE SEPARATION TRANSFORM

The eigenspace separation (ES) transform[11] allows one to reduce the neural-network size while enhancing its generalization accuracy as a binary classifier. The transform projects each training or test pattern into an orthogonal subspace and thereby reduces the dimensionality of the input vector applied to the neural network. The generalization improvement is possible because the orthogonal projection transforms the two classes of patterns in such a way that the vectors of one class tend to have an average length greater than those of the other class. The KL transform preserves information about the inputs in the sense of minimizing the mean square error. The ES transform preserves and makes more accessible the information that is most relevant to a classifier: information that allows separation of data clusters.

The calculation of the transformation matrix U proceeds as follows.

1. Compute the $k \times k$ matrix,
$$\hat{M} = \frac{1}{N_1}\sum_{i=1}^{N_1} X_{1i}X_{1i}^T - \frac{1}{N_2}\sum_{i=1}^{N_2} X_{2i}X_{2i}^T, \tag{2}$$
where X_{1i} is mine pattern i, N_1 is the number of mine patterns, X_{2i} is background pattern i, and N_2 is the number of background patterns.
2. Calculate the eigenvalues of $\hat{M}$, λ_i, $i = 1, 2, \ldots, k$.
3. Calculate the sum of the positive eigenvalues
$$E_p = \sum_{\substack{i=1 \\ \lambda_i > 0}}^{k} \lambda_i, \tag{3}$$
and the sum of the absolute values of the negative eigenvalues
$$E_n = \sum_{\substack{i=1 \\ \lambda_i \le 0}}^{k} |\lambda_i|. \tag{4}$$
4. (a) If $E_p > E_n$, then calculate all m orthonormal eigenvectors that are associated with the positive eigenvalues of $\hat{M}$. Use these eigenvectors as the columns of the $k \times m$ matrix U.

(b) If $E_n > E_p$, then calculate all m orthonormal eigenvectors that are associated with the negative eigenvalues of $\hat{M}$. Use these eigenvectors as the columns of the $k \times m$ matrix U.

(c) If $E_p = E_n$, then either subset of orthonormal eigenvectors can be used in constructing U. In general, choose the smaller subset if possible.

After the matrix U is constructed, $Y_i = U^T X_i$ is computed for each pattern X_i in the training and test data sets. Then each Y_i is applied to the neural network.

9. NEURAL NETWORK ARCHITECTURE AND RESULTS

This section summarizes the performance achieved using a backpropagation neural network for land mine detection. The effect of modifying some of the variable factors in this system are discussed. Some of the relevant result tables are given here and others are available in reference 12.

Several different types of network architecture were investigated. These include: fully connected feedforward networks, partially connected or bottlenecked networks, and recurrent networks. The recurrent networks were found to be too computationally intensive during training to be very useful.[5] The bottlenecked networks were studied in an attempt to reduce the number of weights necessary in the network. However, the generalized KL and ES transforms provided superior methods. Therefore, a fully connected feedforward neural network was considered the best choice for this work.

In addition to the interconnectivity, the size of the network was also studied. It was found that a three-layer (two hidden layers of neurons and one output layer) network gave the best results. A two-layer network had difficulty modeling the complex relationship between the input signals and the classification output, and a four-layer network was unnecessarily complex. Through simulation, the optimum number of neurons for each layer was established. Networks with eight neurons in the first hidden layer, three neurons in the second hidden layer, and one neuron in the output layer consistently showed good results.

From the 88 different measurements taken at each spatial location in the mine lane, 28 components were determined to be the most useful and important. These were the magnitude components of all four S-parameters obtained at the seven frequencies from 880 MHz to 1120 MHz in 40 MHz intervals.

For purposes of training and testing the neural network, the mine lane was divided into five approximately equally spaced sections. The data from the first, middle, and last sections were used for training, and the second and fourth sections were used only for testing, resulting in 7414 training patterns and 5886 test patterns. Because the large number of test patterns allowed an accurate measurement of the network's generalization ability, only the test region mine-detection statistics are used to judge the network's performance.

Preprocessing of the input data before using a neural network proved to be vital for optimum performance. The removal of the underlying "ramp" and "jumps" in the raw data dramatically increased the performance of mine detection. Furthermore, the deterministic removal of outliers from the training data set and balancing the mine data to background data ratio proved to be very beneficial. Both the four-nearest-neighbor and eight-nearest-neighbor outlier removal methods work very well.

Both of the two transform methods used in this study, the generalized KL transform and the ES transform, showed good results. The generalized KL transform with 16 inputs achieved the best peak performance of 100% mine detection with only 10.2% false-positives. However, the ES transform, which used only 16 inputs, was able to produce better average performance than the other methods. Some of these results, using the four-nearest-neighbor outlier removal method, are given in Table 1. Averaged results are given for the training and test data sets using 20 different initial weights and the peak performance obtained. Corresponding results obtained using the eight-nearest-neighbor outlier removal method are given in Table 2.

Table 1: Mine detection performance comparing transform preprocessing methods using the 4-nearest-neighbor outlier removal method.

Transform	Averaged Results (20sims)				Peak Performance			
	Training		Testing		Training		Testing	
Method	% TP	% FP	% TP	% FP	% TP	% FP	% TP	% FP
None [28/28]	99.2	10.4	96.5	12.2	100.0	8.7	100.0	11.4
KL [28/28]	99.4	8.8	95.9	11.2	100.0	8.7	97.9	9.8
KL [16/28]	98.4	11.3	96.0	13.2	100.0	8.5	100.0	10.2
ES [16/28]	99.3	11.1	97.5	12.2	98.5	10.9	100.0	11.7

Table 2: Mine detection performance comparing transform preprocessing methods using the 8-nearest-neighbor outlier removal method.

Transform	Averaged Results (20sims)				Peak Performance			
	Training		Testing		Training		Testing	
Method	% TP	% FP	% TP	% FP	% TP	% FP	% TP	% FP
None [28/28]	99.4	10.5	97.7	12.9	100.0	8.4	100.0	9.8
KL [28/28]	99.2	9.8	98.1	12.9	98.5	7.5	100.0	9.0
KL [16/28]	98.2	11.4	97.6	13.9	98.5	10.8	100.0	15.7
ES [16/28]	98.8	11.4	97.8	13.2	100.0	8.2	100.0	9.3

When evaluating the performance of a classification system, there are two important criteria: the true-positive rate and the false-positive rate. In addition, these two rates are not independent of each other. In fact, there is a strong correlation between the two measures and a tradeoff must be made when maximizing the performance of the system. For the results reported in Tables 1 and 2, emphasis was placed on maximizing the mine detection (true-positive, TP) rate. In most cases, the peak performance was perfect (100% detection was achieved in the test region). Although this best possible detection rate was achieved, there are cases when the false-positive (FP) rate may be undesirably high. Therefore, the question arises whether some of the detection performance can be sacrificed for a better false-positive rate. Some results are presented in Table 3 to illustrate this concept.

In the non-transformed case using the maximum detection criterion, the peak performance achieved was 100.0% TP with 11.4% FP. An averaged measure over 20 different initial conditions resulted in 96.5% TP and

Table 3: Mine detection and false-positive tradeoffs. (non-transformed)

Performance	Averaged Results (20sims)				Peak Performance			
	Training		Testing		Training		Testing	
Criteria	% TP	% FP	% TP	% FP	% TP	% FP	% TP	% FP
Maximum Detection	99.2	10.4	96.5	12.2	100.0	8.7	100.0	11.4
95% Detection	99.8	8.4	95.8	10.2	100.0	8.8	95.8	8.0
90% Detection	99.2	7.2	92.0	8.3	98.5	6.8	91.7	6.3

12.2% FP. However, if we use a minimum detection rate of 90% or 95% as the performance criteria[‡], the results change slightly. For both the best and averaged results, the mine detection rate and the false-positive rate decrease. This demonstrates that the tradeoff between the true-positive detection rate and the false-positive rate is complex. It has not been determined what kind of performance level is desirable or acceptable for practical use. This may be best left to the user of the system.

10. FUTURE METHODS OF MINE DETECTION

We feel that the point of diminishing returns has been reached using the methodology presented thus far. In order to proceed beyond this point, it will be necessary to discover a fundamentally new technique in either the data-gathering or analysis processes. We propose a novel data gathering technique here.

A significant enhancement of the data-gathering system will be the specification of another *independent* feature to add to the collected data vector. This is important because, by adding an independent dimension to the feature vector the network is trained with, a pattern classification problem that was previously not separable may become separable, and hence be completely solvable by the network. After data has been gathered with the new system, we hope to use the existing well developed and studied neural network techniques to realize an even more dependable and effective method of detecting buried land mines.

The method under consideration adds an acoustic sensor such that both acoustic and microwave energy are used to sense underground land mines. A strong acoustic signal is radiated by a loudspeaker located near the surface of the earth, alongside a pair of microwave horns which transmit and receive RF signals. The mine will resonate as a result of the acoustic field imposed on it, causing the surrounding soil to vibrate. These vibrations will be affected strongly by discontinuities, especially hollow spots, like those present in some types of land mines, and voids in the earth caused by mine implantation. The motion of the mine and the earth will cause a modulation of the RF path, resulting in small changes in multipath cancelations. Consequently, the signal detected by the receiving horn will be modulated in amplitude in accord with the level and frequency of vibration. In addition, the RF received signal will be frequency/phase modulated due to both Doppler and multipath effects. The effects of both amplitude and frequency modulation have potential for improving mine discrimination.

11. CONCLUSIONS

The work conducted has shown that a neural network, trained with broadband S-parameter data is able to detect buried anti-personnel land mines with very high accuracy. Furthermore, by varying the design threshold of the rate of desired true positives, the level of false alarms may be made arbitrarily low. Data preprocessing has turned out to be a very important step in the detection process. By "flattening" the data to remove irregularities produced by the measurement process, a very useful data set was obtained. From that point, either the KL or ES transformation was used on the data to reduce the dimensionality of the input data while rotating the underlying coordinate system such that a small network was more easily able to separate the categories of mine and background patterns. Data balancing and the removal of statistical outliers using the four-nearest-neighbor method further aided the network in its classification ability.

Future directions in this work include the addition of an acoustic instrument to the measurement system, which would modulate the microwaves with acoustic energy near the mines' resonant frequency, and would provide an additional independent measurement at each spatial location, thereby reducing the uncertainty of the network's output.

[‡]For these new performance criteria, all neural network results with at least a 90% or 95% TP rate were examined and the one with the lowest FP rate was chosen as the best result.

12. ACKNOWLEDGMENTS

This work was financially supported by ARMY contract DAAK70-92-0003. We would also like to thank Dr. B. Widrow for his constant support and creative input.

13. REFERENCES

1. D.A. Bloom. Land mines keep wars from ever coming to an end. *The Christian Science Monitor*, 87(57):19, February 1995.

2. D. Webster. One leg, one life at a time. *The New York Times Magazine*, January 23, 1994.

3. M.R. Azimi-Sadjadi, D.E. Poole, S. Sheedvash, K.D. Sherbondy, and S.A. Stricker. Detection and classification of buried dielectric anomalies using a separated aperture sensor and a neural network discriminator. *IEEE Transactions on Instrumentation and Measurement*, 41:137–143, 1992.

4. M.R. Azimi-Sadjadi and S.A. Stricker. Detection and classification of buried dielectric anomalies using neural networks—further results. *IEEE Transactions on Instrumentation and Measurement*, 43:34–39, 1994.

5. M.A. Lehr. Adaptive multisource decision-making: Detecting land mines with neural networks using separated aperture sensor data collected at fort belvoir. Final Report, U.S. Army Contract DAAK70-89-K-0001, Stanford University, April 1992.

6. L.S. Riggs and C.A. Amazeen. Research efforts with the waveguide beyond cutoff or separated aperture dielectric anomaly detection scheme. Internal Report #2497, U.S. Army Ft. Belvoir RD&E Center, Ft. Belvoir, VA, 1989.

7. F.R. Clague. Summary of experiments with the separated aperture technique of dielectric anomaly detection, final report. Report number SR-723-29-89, U.S. Army Ft. Belvoir RD&E Center, Ft. Belvoir, VA, 1989.

8. B. Widrow and M.A. Lehr. 30 years of adaptive neural networks: Perceptron, madaline, and backpropagation. *Proceedings of the IEEE*, 78(9):1415–1442, September 1990.

9. R. Anand, K. Mehrotra, C. Mohan, and S. Ranka. An improved algorithm for neural network classification of imbalanced training sets. *IEEE Trans. Neural Networks*, 4(6):962–969, November 1993.

10. R. Schalkoff. *Pattern Recognition: Statistical, Structural and Neural Approaches.* Wiley, New York, 1992.

11. D. Torrieri. A linear transform that simplifies and improves neural network classifiers. Submitted for publication, 1996.

12. T. Doi and G.L. Plett. Adaptive neural networks for mine detection, scientific and technical report. Final Report, U.S. Army Contract DAAK70-92-K-0003, Stanford University, 1995.