

AD A245 006

Semiannual Technical Summary

Artificial Neural Networks for Seismic Data Interpretation

JAN 21 1992

92-01612




30 June 1991

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS




**BEST
AVAILABLE COPY**

Prepared for the Department of the Air Force
under Contract F19628-90-C-0002

Approved for public release; distribution is unlimited.

92 1 17 080

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. The work was sponsored by the Department of the Air Force under Contract F19628-90C-0012.

This report may be reproduced to satisfy needs of U.S. Government agencies.

The ESD Public Affairs Office has reviewed this report, and it is available to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication
FOR THE COMMANDER

Hugh A. Southall

Hugh A. Southall, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission is given to destroy this document
when it is no longer needed.

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY**

**ARTIFICIAL NEURAL NETWORKS FOR
SEISMIC DATA INTERPRETATION**

SEMIANNUAL TECHNICAL SUMMARY

1 DECEMBER 1990 - 30 JUNE 1991

ISSUED 4 DECEMBER 1991

Approved for public release; distribution is unlimited.

LEXINGTON

MASSACHUSETTS

ABSTRACT

This is the second Semiannual Technical Summary of the MIT Lincoln Laboratory Artificial Neural Networks for Seismic Data Interpretation project. The effort during this period has concentrated upon phase labeling and event recognition networks for use in the DARPA/NMRO Intelligent Monitoring System (IMS) for seismic surveillance. Perceptron networks were developed with standard IMS parameters as inputs, improving upon the phase labeling performance of the expert IMS. These networks were developed and tested using data sets containing 5,000 to 10,000 arrivals. An initial version of the expert system achieved a 79% success rate compared with 86% for the neural network. A more recent and improved version of the expert system achieved rates of 87% compared with 90% for an updated neural network. Phase labeling experiments with sonograms and three-component autoregressive modeling for signal representation did not lead to further improvement. Two event labeling experiments were carried out using three-component autoregressive signal models with Radial Basis Function classification networks and involved on the order of 200 events. Success rates were 96.6% for an earthquake/nonearthquake classification experiment and 91% for a Kola Peninsula event recognition experiment. These results, using only autoregressive waveform representations, are encouraging but very preliminary.



Accession For	
NTIS GRANT	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

TABLE OF CONTENTS

Abstract	iii
List of Tables	vii
1. INTRODUCTION	1
2. PHASE LABELING	3
2.1 Phase Labeling Using Standard Parametric Data	3
2.2 Phase Labeling Using Nonstandard Waveform Representations	10
3. EVENT LABELING	15
3.1 Earthquakes versus Nonearthquakes	16
3.2 Kola Peninsula Mines	18
REFERENCES	21

LIST OF TABLES

Table No.		Page
1	Distinguishing Pn from Pg	5
2	Distinguishing Sn from Lg	6
3	Distinguishing Rg from Lg	7
4	Phase Identification Performance (% Correct)	9
5	Performance at Different Arrays (% Correct)	10
6	Phase ID With Autoregressive Modeling	11
7	Phase ID Using Sonograms	13
8	Four-Way Event Identification	16
9	Two-Way Event Identification	17
10	Label Confidences for Error Events	17
11	Single Three-Component Station Event ID	18
12	Kola Peninsula Event Identification (No Dropouts)	19
13	Kola Peninsula Event Labeling (With Data Gaps)	20

1. INTRODUCTION

The overall objective of this research is to develop artificial neural networks to perform specialized recognition functions for nuclear test monitoring systems such as the DARPA Intelligent Monitoring System (IMS) [1]. A specific goal is to develop networks that can be integrated into the IMS to improve its automatic performance.

During this reporting period the greatest amount of work was done on phase labeling, with encouraging results. Initial experiments with event labeling have also been promising. Preliminary arrival time experiments, reported in the previous Semiannual Technical Summary [2], did not lead to promising results with operational IMS data. Specifically, it was not possible to predict with much confidence which IMS arrival time estimates should be corrected, much less predict how much correction should be applied. Furthermore, our discussions with seismologists suggested that event classification aids would be a more significant contribution than improvements in picking arrival time. Thus it was decided to put aside the arrival time estimation work to concentrate on phase and event classification.

Current research has used parametric data from the Center for Seismic Studies (CSS) and an ARCESS/NORESS waveform database of over 200 events that was provided on optical disk by the Science Applications International Corporation (SAIC). Some 50 of the SAIC-supplied events were on hand at the start of this reporting period, and the remainder were received shortly thereafter and integrated into the research database.

Chapter 2 summarizes the work on phase labeling, and Chapter 3 describes some initial event labeling experiments.

2. PHASE LABELING

Phase labeling consists of attaching a phase label (e.g., Pn, Pg) to a seismic arrival. The information used to do this can be standard signal parameters computed from waveforms, contextual parameters such as the arrival order and time delay between arrivals, or nonstandard signal features extracted from the waveforms. Each option was tested, and the best results thus far use only standard signal parameter and contextual parameter information. Detailed results obtained with standard parametric data and nonstandard signal representations follow.

2.1 Phase Labeling Using Standard Parametric Data

During the normal operation of the expert IMS, many parameters are computed to aid in identifying phases and locating events. Because these parameters are readily available and are believed to be useful in the phase identification (ID) process, it seems logical to use them as inputs to a neural network for phase ID.

The process began with simple binary labeling experiments such as using a neural net to distinguish Pn from Pg. These experiments helped to identify the types of parameters and contextual information that are important in the phase ID process. A larger network was then designed to distinguish among all the regional phases. Preliminary experiments and the most recent performance estimates are described here.

Parametric data were obtained from the on-line database at the CSS [3] using Sequential Query Language [4] and the following selection criteria:

- Arrival time within a specified interval
- Associated with some event by expert system and by analyst
- Classified as Pn, Pg, Px, Sn, Lg, Sx, or Rg by analyst
- Valid data available for all parameters of interest.

The arrival time interval varied with the experiment. The initial data set consisted of ARCESS and NORESS arrivals for the time interval January through October 1990. Subsequent experimentation required more than two arrays, and data for the interval November 1990 through April 1991 were obtained when two additional arrays (FINESA and GERESS) were operational. These data were also used for Rg ID experiments, which were added to the expert system at the start of this interval. Data for the interval April through June 1991 have been obtained for the most up-to-date performance statistics.

The neural network used to produce the results described in this section is a simple single-layer perceptron network [5]; multilayer perceptrons were also tried, but it was found that additional layers did not improve performance.

2.1.1 Distinguishing Pn from Pg

To investigate the effectiveness of neural networks on a simple problem, we first experimented with separating Pn from Pg using a small database of 350 arrivals of each type, chosen randomly from the available data. The expert IMS correctly identified 80% of these phases, basing its decision on the following parameters:

- S-P time difference
- P-P time difference
- Number of P phases in this event
- Relative size of P phases in this event
- P-polarization (rectilinearity, emergence angle).

A list was made of the parameters that might be useful to distinguish Pn from Pg, and each was tested individually as an input to a neural network. Because the networks were simple single-layer perceptrons, this is equivalent to putting a threshold on a single input and labeling all phases above the threshold as Pn and all phases below as Pg. After finding the best single input, two inputs were tried and so on, repeating this procedure until the results did not improve by adding more inputs. The results for one and two parameters are shown in Table 1.

The detecting beam parameter in Table 1 is a single bit indicating whether the phase was detected on a horizontal or vertical beam. Slowness and period are calculated by the IMS array processing and spectral analysis software. Rectilinearity, emergence angle, short-axis angle, horizontal/vertical ratio, and maximum/minimum ratio are properties of three-component analysis and are also computed by the IMS signal processing software. The P-polarization parameter is a combination of rectilinearity and emergence angle as used by the expert system. Items such as "only P-type?" are binary parameters that are set to 1 if this phase is believed to be the only P-type phase in the event and set to -1 otherwise. The S-P time difference is the difference between the arrival time of the largest S-phase and the first P-phase associated with the event.

Table 1 clearly shows that the single best parameter for distinguishing Pn from Pg is the S-P time difference. This is physically reasonable because S-P is a good measure of distance, and a P-phase is more likely to be Pg for a nearby event and Pn for one that is distant. When two parameters are used, the best second choice is the bit indicating which P-arrival came first. This works because when both Pn and Pg are present, the first arrival is more likely to be Pn and the second Pg.

When all the parameters on the list are used together as inputs to a neural net, the system achieves about 92% accuracy; the same accuracy could be achieved with only five parameters: S-P, first P, slowness, P-polarization, and max/min ratio.

TABLE 1
Distinguishing Pn from Pg

Input	% Correct	
	Alone	With S-P
Detecting beam	52	83
Slowness	67	81
Period	68	82
Rectilinearity	54	82
Emergence angle	60	82
Short-axis angle	52	82
Horiz/vert ratio	60	82
Max/min ratio	55	83
P-polarization	51	83
Only P-type?	51	81
First P-type?	57	90
Largest P-type?	53	82
S-P time difference	82	

2.1.2 Distinguishing Sn from Lg

Because the results obtained with Pn and Pg phases looked promising, the problem of distinguishing Sn from Lg was also investigated. A small database was created of 350 Sn and 350 Lg phases, chosen randomly from the available data. The expert IMS correctly identified 65%, basing its decision on the following parameters:

- S-P time difference
- Predicted Lg time
- Detecting beam type (horizontal or vertical)
- Relative order of S-phases in this event
- Relative size of S-phases in this event
- S-polarization (short-axis angle, horizontal/vertical ratio).

As before, experiments began by listing parameters that might be useful to distinguish Sn from Lg and testing each as an input to a neural network. The results that were obtained using

three parameters, including the two that gave the best performance by themselves, are listed in Table 2.

The parameters in Table 2 are the same as those in Table 1 with the exception of an added binary parameter, last S-type. The last-S bit is 1 if the phase is the last or the only S-phase in the event. From the table it is evident that last-S is the single most useful parameter in distinguishing Sn from Lg, because often if there is only one S-phase, it is Lg, and if there are two S-phases, the last is Lg. Two parameters (last-S and S-P time) were chosen to be the basis for further experiments because the difference between results obtained with various single inputs was small. Thus the second column represents results obtained with three inputs to the neural network; the best was 86%.

TABLE 2
Distinguishing Sn from Lg

Input	% Correct	
	Alone	With S-P, Last-S
Detecting beam	62	83
Slowness	71	85
Period	53	84
Rectilinearity	62	84
Emergence angle	54	82
Short-axis angle	67	83
Horiz/vert ratio	61	81
Max/min ratio	61	82
S-polarization	65	84
Only S-type?	65	86
O First S-type?	61	86
Last S-type?	76	
Largest S-type?	67	85
S-P time difference	71	

When all the parameters are used as inputs to a neural net, the system achieves about 89% accuracy, which could also be achieved with seven parameters: last S, S-P, first S, slowness, detecting beam, short-axis angle, and rectilinearity.

2.1.3 Distinguishing Rg from Lg

Another phase ID problem is distinguishing Rg from Lg. The expert system does this strictly on the basis of period (or frequency). A simple experiment was performed to determine if other parameters might also be helpful in this task. As before, an equalized database was created that contained (all the available) 218 Rg phases and 218 randomly chosen Lg phases. As the results that are summarized in Table 3 show, period is the single best parameter to distinguish Rg from Lg. When two parameters are used the second should be largest-S, giving a performance of 88% correct. Yet when all the parameters are used, performance is still 88% correct; there is no advantage to using more than the period and largest-S.

TABLE 3
Distinguishing Rg from Lg

Input	% Correct	
	Alone	With Period
Detecting beam	56	86
Slowness	68	86
Period	86	—
Rectilinearity	70	86
Planarity	61	86
Emergence angle	71	86
Short-axis angle	55	86
Horiz/vert ratio	77	86
Max/min ratio	58	86
Only S?	79	86
First S?	83	87
Last S?	54	86
Largest S?	76	88
S-P time difference	63	86

2.1.4 Multiple Category Phase Labeling (Pn, Pg, Px, Sn, Lg, Sx, and Rg)

After the pairwise experiments, testing continued with the more complete set of labels used by the expert IMS and with a larger database. A single-layer perceptron network was used with seven output nodes, one for each phase label. It was discovered that performance could be further improved by adding parameters beyond those used during the binary classification experiments, because they did not include labeling phases such as Px or Sx and, therefore, did not identify useful parameters for that purpose. For example, one added parameter (time delay) measures the difference in arrival times between a phase and one of the same generic type (P or S) that immediately precedes it, expressed as a ratio to the S-P time difference. This parameter is occasionally useful because the analyst is more likely to label as Px or Sx those phases that arrive immediately after a preceding phase.

The network input parameters that were used to obtain the results presented in this section follow in approximate order of importance, i.e., the items at the top contributed significantly to system performance, whereas the items at the bottom affected performance by 0.1% or less.

- Expert system phase
- S-P time difference
- First P/S
- Last P/S
- Period
- Largest P/S
- Slowness
- Time delay
- Rectilinearity
- Short-axis angle
- Maximum/minimum ratio
- Detecting beam type
- Only P/S
- Emergence angle
- Horizontal/vertical ratio.

The final expert system phase (established during multiple array network processing by the expert system) was initially used as the expert system phase parameter. More recent experiments have shown that the neural network operates equally well with the initial phase ("iphase" in the CSS arrival table), which is declared by the expert system using data from a single array.

Networks for phase labeling have been trained using several different data sets, each taken over a different time interval, as described in Section 2.1. Each data set contained 5,000 to 10,000 arrivals, which were randomly divided into two groups of equal size, with half used for training and half for testing. The groups then exchanged roles, and the results obtained from the two testing runs were averaged to get the final performance statistics as given in Table 4.

TABLE 4
Phase Identification Performance (% Correct)

	January through October 1990	November 1990 through April 1991	April through June 1991
Expert system alone (final ID)	79	84	87
Neural net alone	82	85	87
Neural net with expert system ID input	86	89	90

The neural network improved upon the expert system performance in every case, regardless of the time interval used. As the expert system improved over time, neural network performance improvements became somewhat smaller but still significant in the most recent results. Neural network performance without expert system input is similar to expert system performance.

These results show that a neural network would improve the phase ID performance of the expert IMS. The parameters needed by the neural network are readily available, because they are the expert system's initial phase ID and others that are used by the expert system. The next logical step in this research is to integrate phase labeling networks into the existing expert IMS for refinement and extensive testing. The current plan is to work with SAIC toward this goal.

2.1.5 Performance for Different Arrays

Performance sensitivity to the array used for training is an important issue. Will a network trained on one array give good performance when used with another? To investigate this question, IMS data were used from the time period November 1990 through April 1991 to train four different networks, one for each array. All four networks were then tested on data from all four arrays. (In cases where the training and the testing array were the same, the data were split in half, and the array was trained on one half and tested on the other.) The results are given in Table 5.

TABLE 5
Performance at Different Arrays (% Correct)

Training	Testing			
	ARCESS	FINESA	GERESS	NORESS
ARCESS	87	87	90	89
FINESA	83	90	88	89
GERESS	83	87	90	88
NORESS	85	87	90	89

Performance is best when the training and the testing data are from the same array; however, when half the data from all four arrays is used for training and the other half for testing, performance is very similar to that obtained with training and testing on each array individually. (The average of the diagonal entries of Table 5 is 89%, which is equal to the 89% reported for the same time period in Table 4.) Tentatively, the conclusion is that these four arrays are similar enough that separate neural networks may not be necessary for each. Additional experiments with more data will be required to confirm this result.

2.2 Phase Labeling Using Nonstandard Waveform Representations

Experiments have also been conducted with alternative waveform representations to see if useful phase labeling features could be uncovered that are not routinely extracted by the IMS signal processing software. The objective was to develop networks that could either perform the phase labeling themselves or provide additional useful inputs to other networks such as those discussed earlier. A key problem is finding a waveform representation that captures essential discrimination features using only a modest number of parameters.

Two waveform representations, autoregressive parameters and sonograms, were tested. Neither approach by itself improved upon the performance of the expert IMS; however, it may be possible to combine the parameters used in Section 2.1 and the new parameters derived from a waveform representation to create a new neural network with better performance than either one alone.

Waveform experiments used subsets of the waveform database provided by SAIC, which contained a few hundred events. All the data were from the September 1989 through October 1990 interval. For the subsets that were used, the expert IMS phase labeling success rate was in the 76% to 81% range, which compares with the 79% expert IMS success rate reported in Section 2.1.4 for the same time period for a much larger number of arrivals.

2.2.1 Autoregressive Representation

One way to represent waveform data is to use an autoregressive model [6], which is common in many signal processing applications and occasionally used for seismic signals [7] but is not standard. Most applications have involved single waveforms—the output from a microphone for speech applications and the output of a single, usually vertical, seismometer for seismic applications. A decision was made to investigate the use of autoregressive models of three-component seismic signals for phase labeling in the hope that such models might capture useful frequency and relative phase information between the components that were missed by standard IMS signal measurements.

The experiments used three-component beamformed data derived from the waveform database supplied by SAIC. Waveforms were segmented into nonoverlapping time windows, using the phase arrival times provided by IMS. Each time window was considered individually, that is, completely out of context of any other arrivals in the event. The horizontal components were rotated into radial/transverse coordinates using the azimuth estimate provided by IMS. Slowness and direction used for beamforming were also provided by IMS. The three-component beams were processed using fourth-order, multichannel, autoregressive analysis [6]. The calculated parameters for each time window then became the inputs to a neural network classifier.

The slowness of the arrival was also used to make a preliminary discrimination into P-class, S-class, teleseisms, and noise. A supervised neural network classifier called a Radial Basis Function (RBF) neural network [8] was used to discriminate between members of the same class, for instance, Pn from Pg from Px. The input to the RBF consisted of the parameters calculated from the autoregressive analysis.

Results using the RBF network are summarized in Table 6, with the first line indicating results obtained with arrivals that are separated by enough time so that only one occurs in the analysis time window, thus the autoregressive parameters are not influenced by the presence of other arrivals. The second line indicates results with all available data, regardless of whether additional arrivals occurred in the analysis window.

TABLE 6
Phase ID With Autoregressive Modeling

Experiment	Results (% Correct)
Single arrivals	75
Possible multiple arrivals	63
Note: Expert system performance was 80.5%.	

The results with autoregressive modeling are close to those of the expert system for the single arrival cases. When the time between arrivals is short, however, this approach has difficulty differentiating between multiple arrivals within the analysis window, and other information becomes necessary. It may be possible to improve network performance by adding contextual information to the RBF inputs or by combining the autoregressive inputs with the parameter-based network that was discussed earlier.

2.2.2 Sonogram Representation

An alternative to the autoregressive model is a sonogram representation. The motivation for using sonograms is to exploit the time-varying frequency content of seismograms by teaching neural network classifiers the time-varying frequency patterns common to different phases. Sonograms were computed with a short-time Fourier transform [9] with overlapping Welch windows [10] to ensure gradual transitions in both the time and frequency domains.

The neural network classifier used for this work is based on the ART2 classifier [11]. One reason for using ART2 was to gain more experience with its unsupervised learning mode, which makes it a candidate for future work involving the detection of novel events. ART2 was extended by the addition of an extra layer to provide supervised labeling of its categories. The weights connecting the ART2 output nodes to label nodes are intended to represent the probabilities that the labels are correct. The resulting overall structure is referred to as a "semisupervised ART2 network."

The data used in these experiments were segments of unfiltered detecting beams included in the databases supplied by SAIC. Phase arrivals were segmented on the basis of IMS time picks. Segments started 1.6 s before the IMS pick and lasted for a total of 20.4 s. Time windows of 1.5 s with 80% overlap were used to calculate the sonograms. The spectra were normalized using spectra from a 20.4-s segment in the coda of the event, defined to be 2.4 times the difference between the onset of the largest S and the first P phases (both determined by IMS) in this event. The motivation for the normalization is to mitigate the effects of background noise, hide the spectral modulation that occurs in mining blasts [12], and equalize the instrument response across all frequencies. After normalization the data are logarithmically sampled along the time axis so that the onset of the phase is most heavily weighted. The result is a two-dimensional sonogram that is raster-scanned and used as input to one of three enhanced ART2 networks of the type described. The determination as to which network to use is based on slowness, which is provided by the IMS. The three networks represent P-class, S-class, and teleseisms or noise. According to the results in Table 7 the performance of the expert IMS is superior to that of the sonogram/ART2 labeler. It remains to be determined if sonogram/ART2 outputs can be combined with other parameters to improve performance.

TABLE 7
Phase ID Using Sonograms

Experiment	Expert System (% Correct)	Neural Network (% Correct)
ARCESS	76	68
NORESS	80	66

3. EVENT LABELING

The goal of the current event labeling research is to develop analysis aids that can be integrated into the IMS. These aids might flag events for more careful analyst review and make the processing of typical events more routine. The hope is that networks might identify certain mines, quarries, or seismic areas on the basis of typical waveforms. An event flagged as atypical would be a candidate for more analysis, independent of the results of other processing elements in the IMS. Disagreements between networks and those other elements would also be a warning that more analysis is necessary.

Two exploratory event labeling experiments have been performed. One, an earthquake/non-earthquake classification experiment, achieved a 96.6% success rate and provided an example of how networks might flag unusual events for more detailed analysis. Another experiment was performed in which an attempt was made to recognize Kola Peninsula mines without using azimuth or any other parametric information from the expert system; this experiment achieved a 91% success rate.

Waveform data and event labels provided by SAIC were used for event labeling experiments. SAIC provided *Helsinki Bulletin* classifications as labels when they were available and assigned labels for other events. Three primary labels were used by SAIC: mine, earthquake, and explosion. Explosions are events that are probably not earthquakes but are not located at known mine sites. Many, if not all the explosion events could be mine events. In addition, some events were left unlabeled.

An RBF neural network [8] classifier was used with feature vectors derived from a multichannel autoregressive analysis [6], which was performed on a single time window containing all the associated arrivals for each event, starting at the first arrival. Three component beams, steered to the first P arrival speed, as well as single three-component instruments, were used in the experiments. A minor waveform data problem was discovered; there were gaps in the waveform data files stored on-line at CSS. The gaps, which seem to indicate a transmission problem, are almost always an integral number of seconds and may occur at different times on the various components of the stations within an array. Although not as obvious when beamformed data are used, the gaps are obvious when single stations are processed. Some of the experiments reported below were performed on single three-component stations to be certain that the data gaps were not significantly affecting the results.

Event labeling experiments were performed using leave-one-out testing. Given N events, the classifier was trained on $N - 1$ and tested on those remaining. The entire training/testing procedure was repeated N times to test on all N events, individually. The total number of events was on the order of 200, and this approach maximized the use of this relatively small number.

3.1 Earthquakes versus Nonearthquakes

Three different earthquake labeling experiments were performed using ARCESS waveform data. The first trained networks to reproduce the event labels provided by SAIC. The waveform data used were three-component ARCESS beams steered to the first P phase arrival. Beams with data gaps were excluded but those for which the contributing three-component sensors might have gaps were included (see Table 8). Unfortunately, the data included one event for which SAIC did not provide a label. This produced minor peculiarities because a network cannot be meaningfully trained and tested with a single event. In addition, as noted previously events that are labeled explosions may in fact be mine blasts, and perhaps it is unfair to penalize the network labeler for confusing these two classes. With this in mind, only the two events that the neural network labeled earthquakes but the analyst did not, and the four events that the analyst labeled earthquakes but the neural network did not, might be considered errors. With this interpretation, the neural network misclassified 6 of 178 events for a performance of 96.63% correct.

TABLE 8
Four-Way Event Identification

Analyst Labels	Neural Network Labels				Total	Percent
	Unlabeled	Mine	Earthquake	Explosion		
Unlabeled	0	1	0	0	1	1
Mine	0	137	1	5	143	80
Earthquake	0	3	10	1	14	8
Explosion	0	14	1	5	20	11
Total	0	155	12	11	178	
Percent	0	87	7	6	100	
Total correct was 85.29%.						
6/178 earthquake errors yields 96.63% correct.						

The second experiment was similar, but involved only two categories: earthquakes and non-earthquakes. For this experiment all events that were not labeled as earthquakes by SAIC, including the unlabeled one, were treated as nonearthquakes. Table 9 shows that the distribution of errors was slightly different from the previous experiment, but overall performance was similar with about 97% correct.

TABLE 9
Two-Way Event Identification

Analyst Labels	Neural Network Labels		Total	Percent
	Earthquake	Nonearthquake		
Earthquake	9	5	14	8
Nonearthquake	1	163	164	92
Total	10	168	178	
Percent	6	94	100	
Total correct was 96.63%.				

It is interesting to examine the errors made by the neural network in the dichotomous classification problem. The events for which errors were made are listed in Table 10, which shows the confidences (probabilities) that the neural network assigned to the correctness of its labels for these events. The second event (IMS ORID 191806) stands out because the neural network assigned a confidence of 99.9% to its assessment that this event was not an earthquake. This event was reported in the *Helsinki Bulletin*, but it was not given an identification, which suggests that the analyst was unsure about it. The event occurred at 1:38 A.M., and it was not near any catalogued mines. SAIC assigned its own "earthquake" label, a subjective low confidence measure. This is an example of an unusual event that a neural network might help to identify as one that needs more careful scrutiny.

TABLE 10
Label Confidences for Error Events

Earthquake	Other	ORID	Event Date/Time
0.592	0.408	198311	90 05 23 / 19:04:24.003
0.001	0.999	191806	90 02 04 / 01:38:32.574
0.093	0.907	193280	90 04 02 / 13:46:23.256
0.487	0.513	198272	90 05 20 / 10:27:07.366
0.494	0.506	128464	89 10 17 / 20:42:08.973
0.203	0.797	160940	89 11 10 / 07:06:14.404

The previous results were obtained using three-component, beamformed array data. Beamforming increases the signal-to-noise ratio and reduces the effects of data gaps in the waveform files. On the other hand, because phase velocity is also a parameter needed to compute the beams, a beam cannot be optimized for both S- and P-type phases. To investigate this collection of effects, the previous results were compared with earthquake/nonearthquake labeling using a single three-component station with and without events with data gaps.

Table 11 shows that for individual three-component seismometers, there is not a significant difference with and without the data gaps; however, the 92% to 93% performance with single three-component signals is significantly less than the 96.6% success rate obtained using three-component beams. Apparently, for this experiment the presence of data gaps was less important than the improved signal-to-noise ratio that beamforming provided.

TABLE 11
Single Three-Component Station Event ID

Gaps	Percent Correct	Events
Without data gaps	92.61	71
With data gaps	92.96	230

3.2 Kola Peninsula Mines

The Kola Peninsula is an active mining region; approximately 25% of the events in the database that were received from SAIC were from that region. All the Kola events were labeled as mining explosions by SAIC. The experiments reported in this section involve discrimination between Kola Peninsula events and all others in the current database. For these experiments azimuth and distance were not provided to the neural network, although the horizontal components of the data were rotated to correspond to the approximate radial and transverse directions estimated by the expert system for the first arrival of the event. Like the experiments in Section 3.1, labeling was based solely on the coefficients derived from the autoregressive analysis of the waveforms.

Table 12 shows Kola Peninsula event classification results that were obtained by using data from the central ARCESS three-component station. Only waveforms without data gaps were used. Expert IMS performance was estimated by comparing the automatic expert system locations with the final IMS output after analyst corrections were applied. Neural network performance was

estimated by comparing it with the final IMS output. Using only waveform information, the neural network correctly identified Kola events almost as well as did the expert system when it used explicit array-based event location algorithms. The percent Kola and non-Kola events were about the same for the expert system and the neural network, but they did not make errors on the same events.

TABLE 12
Kola Peninsula Event Identification (No Dropouts)

Analyst-Adjusted Locations	Neural Network Labels		Total	Percent
	Kola Pen	Other		
Kola Pen	10	2	12	18
Other	4	50	54	82
Total	14	52	66	
Percent	21	79	100	
Total correct was 90.91%.				
Estimated expert system performance was 95.45%.				

To gain a better understanding of data gap effects, networks were also trained and tested using only signals containing data gaps with the outcome shown in Table 13. A comparison with Table 12 suggests that the gaps were short enough so that they did not markedly influence the results.

Beams steered to the first P arrival were also used to train and test a Kola Peninsula mining event labeler. The performance, 89% correct, was not as good as for a single three-component station. The reason may be that the beamforming velocity filtering is discarding useful information from other phases. This possibility will be investigated using separate beams steered to each detected and associated phase.

TABLE 13
Kola Peninsula Event Labeling (With Data Gaps)

Analyst Labels	Neural Network Labels		Total	Percent
	Kola Pen	Other		
Kola Pen	32	9	41	19
Other	10	168	178	81
Total	42	177	219	
Percent	19	81	100	
Total correct was 91.32%.				
Estimated expert system performance was 96.80%.				

REFERENCES

1. T.C. Bache, et al., "Intelligent Array System: Final Technical Report," San Diego, Calif.: Science Applications International Corporation Technical Report (1990).
2. "Artificial neural networks for seismic data interpretation," Lexington, Mass.: MIT Lincoln Laboratory, Semiannual Technical Summary (30 November 1990).
3. J. Anderson, W.E. Farrell, K. Garcia, J. Given, and H. Swanger, "Center for Seismic Studies Version 3 Database: Schema Reference Manual," Center for Seismic Studies Technical Report C90-1 (1990).
4. J. Anderson and H. Swanger, "Center for Seismic Studies Version 3 Database: SQL Tutorial," Center for Seismic Studies Technical Report C90-02 (1990).
5. R.P. Lippman, "Pattern classification using neural networks," *IEEE Commun. Mag.*, 47-64 (1989).
6. R.J. Jones, "Prediction of multivariate time series," *J. Appl. Meteorol.* **3**, 285-289 (1964).
7. O. Sandvin and D. Tjostheim, "Multivariate autoregressive representation of seismic P-wave signals with application to short-period discrimination," *Bull. Seismol. Soc. America* **68**, 735-756 (1978).
8. D.S. Broomhead and D. Lowe, "Multi-variable functional interpolation and adaptive networks," *Complex Systems* **2** (1988) 312-355.
9. R.A. Roberts and C.T. Mullis, *Digital Signal Processing*, Reading, Mass.: Addison-Wesley, (1987).
10. W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, *Numerical Recipes in C*, New York: Cambridge University Press (1988).
11. G.A. Carpenter and S. Grossberg, "ART2: Self-organization of stable category recognition codes for analog input patterns," *Appl. Opt.* **26**, 4919-4930 (1987).
12. M.A.H. Hedlin, J.B. Minster, and J.A. Orcutt, "The time-frequency characteristics of quarry blasts and calibration explosions recorded in Kazakhstan, USSR," *Int. J. Geophysics* **99**, 109-121 (1989).

