ESANN'2005 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 27-29 April 2005, d-side publi., ISBN 2-930307-05-6.

Performance of EMI Based Mine Detection using Back-Propagation Neural Networks

Matthew Draper and Taskin Kocak

Dept. of Electrical and Computer Engineering
University of Central Florida
Orlando, FL 32816
mdraper@mail.ucf.edu, tkocak@cpe.ucf.edu

Abstract. We propose and evaluate a neural network approach to mine detection using Electromagnetic Induction (EMI) sensors which provides a robust non-parametric approach. In our approach, a neural network with the well-known back-propagation learning algorithm combines the S-Statistic with the δ-Technique to discriminate between non-mine patterns and mines. Experimental results show that this approach reduces false alarms substantially over using just the δ-Technique or the energy detector.

Key-words: Mine Detection, Back-Propagation, δ-Technique, S-Statistic, False Alarm Filtering.

1 Introduction

Antipersonnel landmines are devastating weapons of war, but they are equally devastating after a war. The vast majority of landmines in use today around the world have no means of self-neutralization or self-destruction. Millions of anti-personnel mines are estimated to be buried in the ground of forty countries. They kill or maim more than 2000 civilians per month and prevent the return to productive activities of vast areas of land. Demining activities are supported by several humanitarian organizations, at an estimated cost of $800 per mine found [1]. Therefore there is a real need for technologies which can render demining more effective, more cost-effective and safer.

We propose an artificial neural network (ANN) architecture for post-processing of mine field data, denoted $BPNN(\delta, S)$. This architecture is a feedforward neural network model which is trained using the back-propagation learning algorithm on two features extracted from the data in the mine field: the δ-value and the S-Statistic. This approach uses data from a small calibrated area to train the corresponding neural network, which is then used for mine detection over much larger areas. Our experimental evaluation, using available sensory data [2] shows that the trained network architecture can be effectively used in areas which are geographically remote from the calibration area. It is also effective when tested with sensory data obtained with EMI sensors which have different characteristics from those which were used to collect the network training data.

1.1 The Minefield Data

The minefield data we use in the present study is based on measurements provided by DARPA [2], with two different electromagnetic induction sensor systems, at a variety of geographic locations. These locations are denoted as Firing Point (FP) 20, Firing Point (FP) 22, Seabee, and Turkey Creek.

One source of inaccuracy in the practical use of the data we employ is related to the exact location of the sensor being used as data was registered. This is due to a variety of instrumentation and data collection effects, leading to errors in registering the sensor's position as it travels continuously across the minefield. Hence, we have followed a commonly accepted procedure suggested for using this data, which is to register the mine locations by analyzing the energy levels near the approximate known mine locations, if a location is known to contain a mine, then energy measurement at its immediate neighbors must be lower. To give an idea of this effect, one of the 5m x 5m regions that we examine is shown in Figure 1.

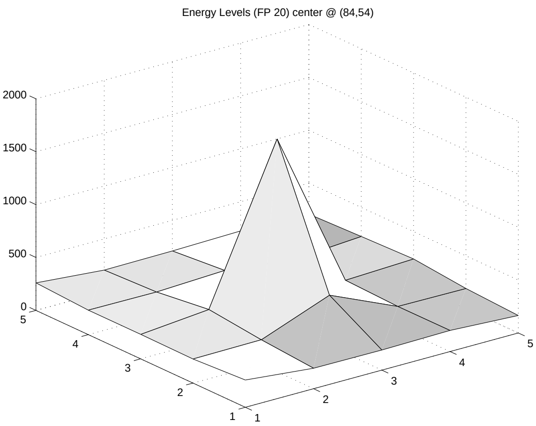

Fig. 1: Energy profile at a mine location in FP20. The z-axis show the energy value of the coil, and x and y axes are relative x and y positions in the minefield.

2 The Energy Detector, the δ-Technique and the S-Statistic

An energy detector is a simple and useful detection technique which will report an "alarm" – i.e. a location which may possibly contain a mine – on the basis of some measured response energy value which exceeds some given threshold [2, 3]. Since it is the simplest possible detection technique, we will use it as a basis for comparison with other methods. With a low enough threshold energy value, an energy detector will yield very high probabilities of detection of mines but will also lead to unacceptably high false alarm rates. Handling false alarms in the minefield can be almost as expensive as removing a real mine. Since thousands of false alarms can occur in sweeping a relatively small minefield, it is important to be able to devise techniques which provide a high probability of detection, with false alarm rates which are much lower than those resulting from the energy detector.

One such improved detection technique is the δ-Technique reported in an earlier paper [4] which significantly reduces false alarm rates by making use of neighborhood or area information around each location. It determines the number of neighboring points (there are 8 of them) that have a higher energy than the current point. If this number is greater than a threshold that you set (usually 7 or 8), then the point is considered to be a landmine. The idea in using this statistic is to stress that relative energy values are more significant indicators of the presence of a target, than absolute levels or differences in energy.

Consider again the energy measurement shown in Figure 1, where we see that the energy at the mine location is higher than that at neighboring points. If we assume that this is generally true for most mine locations, and we further assume that in non-mine locations this property does not hold true, then we would have a very good indicator that will help us identify mine locations. Just as the δ-Technique exploits this property, we propose a new and very effective statistic using this type of local difference information, which we call the $\underline{S}$-Statistic, where:

$$S = \frac{E(p) - (8-m)/8 \sum_{All\ p_n} E(p_n)}{E(p)} \tag{1}$$

where $m = 7$ or $m = 8$.

3 A Neural Network Mine Detector Using the δ-Technique and the S-Statistic

For our approach, we combine the S-Statistic with the δ-Technique in a neural network design. The learning algorithm exploits data from a small calibrated area to train a neural network which can then be used for mine detection over much larger areas. The network is only trained with Z-coil $1m$ data from the $30m \times 15m$ calibration area of a site known as Firing Point 20 (FP20) [2].

We use a three layer feedforward back-propagation neural network (BPNN) to detect mines and reject false alarms. The network has two input neurons. When the network is either trained for some location p, or when it is asked to provide a decision (mine or non-mine) for the location p, one input neuron receives the input s (see Equation 1) and the other receives $\delta = m/8$ where δ is the δ-Technique parameter and m is the number of immediate neighbors required whose energy values are strictly less than the center point's energy value.

In the network's output layer there are two neurons which are used to decide between the two hypotheses (mine or a non-mine) for the location for which input data is presented. The network has six intermediate (hidden) layer neurons.

In the decision phase, when the network is being applied to data it has not observed previously, we use a decision variable (D) which is simply the ratio of output M (hypothesis that the current location is a mine target) to output N (hypothesis that the current location is not a mine). When the input values s and δ for a given location are presented to the network and $D > 1$ the location is declared to contain a mine; otherwise it is declared not to contain a mine. Clearly, just as with any other detection algorithm, the neural network is not "perfect" so that the probability of a false alarm is not 0, and the probability of correct detection of a mine is not 1. However, as we shall see below, its performance is remarkably good with little training and across different EMI sensors.

4 Experimental Results

This section summarizes the performance achieved using neural networks for land-mine detection. After training the proposed neural network architecture on the calibration data, it was tested for all available data which includes measurements from both $1m$ and $0.5m$ EMI sensor systems (8 separate sets), for all the data including calibration and "center square" areas. The center square is a $100m \times 100m$ area in which registration targets are placed. Since the energy measurements vary from one site to another and also for different sensory systems, we prepare the results with zero threshold energy level. We report the results of the energy detector, δ-Technique for $\delta = 7/8$, and the $BPNN(\delta, S)$.

The results for the four minefield sites with two different EMI sensors, are given in Table 1. For all sites we observe that the ANN based technique achieves substantial reduction in the probability of false alarms over the δ-Technique and the energy detector, though it may not find as many actual targets as the δ-Technique.

Table 1: ANN Improvement for Reducing False Alarms for Different Sites with 1m and 0.5m Z-coil Data

Location Names	FP 20		FP 22		Seabee		Turkey Creek	
Data Sensor	1m	0.5m	1m	0.5m	1m	0.5m	1m	0.5m
Points searched	8406	7896	7687	7896	11134	10395	8109	7945
No. of Mines	21	24	24	23	24	24	24	24
FA detected: Energy det.	8385	7872	7663	7873	11110	10371	8085	7921
Mines detected: Energy det.	21	24	24	23	24	24	24	24
FA detected: $\delta = 7/8$	2067	1381	1862	1475	2628	1746	2014	1463
Mines detected: $\delta = 7/8$	21	23	22	22	24	24	24	24
FA detected: $BPNN(\delta, S)$	978	588	868	615	1291	853	985	719
Mines detected: $BPNN(\delta, S)$	20	23	21	21	24	23	24	24

The Receiver Operating Characteristic (ROC) for FP20 is plotted as shown in Figure 2. This curve represents the relation between the probability of detection and the probability of false alarms for a certain detector. Three ROC curves are plotted: (i) for the $BPNN(\delta, S)$, (ii) for the δ-Technique, and (iii) for the pure energy detector.

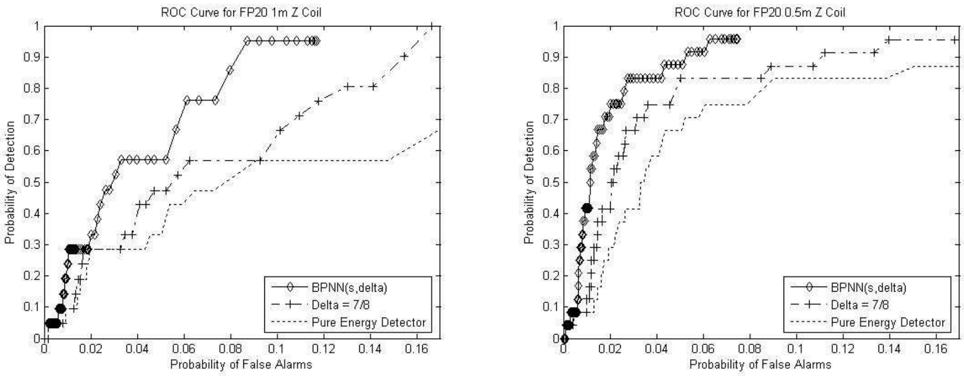

Fig. 2: ROC curves for FP20 (a) 1m Z coil data (b) 0.5m Z coil data

It can be seen that the ANN detector provides better performance than the δ-Technique, and both have significantly better performance than the pure energy detector. For example, we can see from Figure 2 that a 0.08 false alarm probability is obtained when the probability of detection is 0.5 for the pure energy detector, 0.57 for the δ-Technique, and 0.80 for the $BPNN(\delta, S)$. We noticed similar improvements for the data from the other sites.

Of course, the ANN based technique requires training and is, therefore, more complex and computationally more costly than the δ-Technique. Notice also that for a certain value of the probability of detection, there may be multiple values of the percentage of false alarms. This is because it sometimes occurs that, as we vary the energy threshold, the probability of detection remains unchanged while the false alarm probability varies.

5 Conclusions

In this paper it is shown that the proposed network is very effective in detecting mines and rejecting false alarms. From our experimental results we can see that an ANN detector offers a robust non-parametric technique for mine dection which can significantly out-perform the δ-Technique and the energy detector. This detector also has the advantage that it can be trained on limited calibration data and still perform robustly when applied to a variety of geographic locations. Despite high amounts of clutter in the mine field data, this detector outperforms other techniques in determining the difference between clutter and actual mines. Finally, it has been shown in our experiments that the same detector, trained using one EMI sensor, can be used with data gathered from another sensor. This indicates that the ANN approach is robust enough to be used with different sensors which would allow for the mine detector to be forward compatible with future sensor technologies.

References

[1] US Dept. of State, "Hidden killers - the global landmine crisis," Report to the US Congress, 1994.

[2] V. George, T. Altshuler, et al., *Background data collection plan,* DARPA Defense Science Office.

[3] A. Andrews, V. George, and T. Altshuler, "Quantifying performance of mine detectors with fewer than 10,000 targets," in *Detection and Remediation Technologies for Mine and Minelike Targets,* SPIE Proceedings vol. 3079, pp. 273-280, April 21-24, 1997.

[4] E. Gelenbe and T. Koçak, "Area based results for mine detection," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 38, no. 1, pp.12-24, Jan. 2000.