

Classification of Sonar Signals Using Bayesian Networks

Michael J. Larkin
Naval Undersea Warfare Center (Code 3111)
Newport, RI 02841-1708
larkin_m@vsdec.npt.nuwc.navy.mil

Abstract

A pertinent issue in the development of automated classifiers is the combination of information extracted from the data with a priori information that we may have about the classes in question. Incorporating prior knowledge is particularly useful if the data is incomplete or imprecise. Bayesian networks offer an ideal representation for the combination of a priori knowledge with data. This paper discusses the implementation of a Bayesian network classifier for the purpose of classifying underwater mines, using knowledge that is known by experts (sonar operators) about the distinguishing characteristics of mines, such as size, shape, shadow and resonance.

1. Introduction

One of the more pertinent issues in the development of automated classifiers is the combination of information extracted from the data with a priori information that we may have about the classes in question. Theoretically, performance of a classifier should be improved if we are able to incorporate what we know about the objects that we are attempting to classify. Incorporating prior knowledge is particularly useful if the available data is incomplete, imprecise, or noisy; thus not lending itself to creation of a data-driven classifier.

The problem at hand, classification of underwater mines from sonar returns, is a particularly interesting problem, and a good demonstration for this approach. Automated classification of mines, while having been well-studied, is generally difficult in the real world because of factors such as noise in the observations (especially in the shallow water environment), the presence of objects that have similar characteristics to mines, and occlusion by the bottom or by other objects.

Experienced sonar operators know what to look for in the image presented on a sonar scope- they are trained to look for these features and, with experience, are able to

determine their presence in a sonar image and assign a level of confidence to the presence of each feature. The physical features that they observe include size, shape, shadow (a light area on the sonar return that is less insonified than the surrounding region; most often caused by signal blocking from an acoustically opaque object), and resonance (multiple echoes produced by targets such as mines due to acoustic reflections internal to the target.)

Bayesian networks offer an ideal representation for the combination of a priori knowledge with information extracted from the observed data. The subject of this paper is the implementation of a Bayesian network-based classifier for the purpose of classifying objects detected in sonar returns as minelike or non-minelike. By combining as much information as can be extracted from the data with knowledge obtained from experts (sonar operators), classification performance is enhanced.

2. Bayesian networks

A Bayesian network is a graphic model that permits probabilistic relationships among a set of variables of interest to be encoded. Bayesian networks have become a popular method of encoding uncertain expert knowledge into expert systems [1],[2]. More recently, methods have been developed to learn Bayesian networks from data.

Bayesian techniques, when combined with certain statistical techniques, offer an effective method for the analysis of data. Bayesian networks encode dependencies among all the variables, and thus are more robust to incomplete data. This is particularly important in the case where two of the input variables are strongly anti-correlated. Most models will produce inaccurate results when one of the two variables is not observed, but because the correlation is encoded in the Bayesian network, this problem is overcome.

Additionally, this model has both causal and probabilistic semantics and hence is an ideal representation for combining domain knowledge with observed data. Prior knowledge is encoded in a straightforward method,

and the strength of the causal relationships are encoded by assigning probabilities.

Another advantage of Bayesian techniques is an effective approach for avoiding the overfitting of data. In the case where data is meager, there is no necessity to hold out some of the data for testing. The model can be smoothed in such a way to allow all available data to be used for training.

Bayesian networks encode properties of probability distributions using directed acyclic graphs (DAGs). A Bayesian network is a pair (D, P) where D is a DAG and P is a probability distribution called the underlying distribution. Each node i in D corresponds to a variable X_i in P , a set of nodes I corresponds to a set of variables \mathbf{X}_I and x_i, x_I denotes values drawn from the domain of X_i and the cross product domain of \mathbf{X}_I , respectively. Each node in the network is regarded as a storage cell for the distribution $P(x_i | \mathbf{x}_{\pi(i)})$ where $\mathbf{x}_{\pi(i)}$ is a set of variables that correspond to the parent nodes of i . The underlying distribution represented by a Bayesian network is composed via

$$P(x_1, \dots, x_n) = \prod P(x_i | \mathbf{x}_{\pi(i)}) \quad (1)$$

The role of a Bayesian network is to record a state of knowledge P , to provide a means for updating the knowledge as new information is accumulated and to facilitate query answering mechanisms for knowledge retrieval. A standard query for a Bayesian network is to find the posterior distribution of a hypothesis variable X_I , given an evidence set $\mathbf{X}_J = \mathbf{x}_J$, i.e., to compute $P(x_I | \mathbf{x}_J)$ for each value of X_I and for a given combination of values of \mathbf{X}_J . The answer to such queries can, in principle, be computed directly from equation (1) because this equation defines a full probability distribution. However, treating the underlying distribution as a large table rather than as the composition of several small ones might be computationally inefficient, unless we exploit independence relationships encoded in the network. Algorithms that efficiently find the independencies in the topology of a Bayesian network, such as d -separation [3], typically do so by testing for conditional independence.

3. Feature extraction in mine detection

The application problem under consideration is the detection of underwater mines. The intent is to construct an automatic classifier that will render a decision as to whether or not an object detected by a minehunting sonar is in fact a mine. Previous work by the author [4] has emphasized the problem of feature extraction as means of efficiently preprocessing signals and deriving a set of

features to input into a classifier, thus enhancing classification performance. Feature extraction reduces the complexity of a highly dimensional signal, while preserving the important information inherent in its structure. Many of the approaches to feature extraction are analogous to principal component analysis (PCA); one technique that we have successfully used is the BCM algorithm [5].

This work led to the problem of determining the characteristics of the features that the BCM neurons identify in minelike objects. For example, a trained sonar operator determines that an object is minelike by observing several characteristics (size, shape, reverberation, shadow, aspect) in the sonar image. Because we employed a biologically-motivated method to extract features from the input signals, a natural question was to ascertain if the features derived by BCM are analogous.

In the case of mine detection, we note that there exists a defined set of characteristics or features that sonar operators are trained to detect, and this is their basis for determining whether or not a sonar contact is minelike. We thus can construct a classifier which actually consists of a set of classifiers- each one performing a classifier to determine whether or not a specific feature is present in the sample. By combining the results of all the feature classifiers, we can then obtain an overall classification result.

In order to do this, and obtain the probability $P(x = \text{mine})$ for the sample x , we need to know each of the probabilities conditional on the features, for example, $P(\text{mine} | \text{shadow})$, the probability that a contact is mine given that a shadow is present. This probability is not immediately determinable. However, the prior probability $P(\text{shadow} | \text{mine})$ can be obtained from the expert. This probability is determined by the skill of the operator and the characteristics of the sensor. We can also make a determination of the probability $P(\text{mine})$ that any random contact is a mine, based on (operational) intelligence. By simple application of Bayes' Rule, we are thus able to compute $P(\text{mine} | \text{shadow})$, and likewise for the other features. Hence, we can then compute the probability that the contact is a mine.

4. Experimental results

In order to assess the viability of this approach, a multiple layer classifier was constructed, as depicted in Figure 1. In the first layer of the classifier, feature detectors determine the presence of each of the characteristics that will be used for classification. The output of the feature detectors are then weighted, based on the posterior probabilities described above, and fed into a classifier which makes the determination whether or not the object is a mine.

The classifier was tested on a set of sonar data obtained from a test range, with multiple returns from twelve objects (seven dummy mines, and five minelike objects such as rocks). In order to construct the prior probabilities necessary to implement the Bayesian network, personnel trained in mine detection were consulted (the author is a Naval Reserve officer and has spent many hours aboard minesweepers conducting minehunting operations).

To provide some basis of discerning the validity of this approach, a comparison experiment, using the same set of data, was performed using the BCM preprocessor. Here, the feature extraction algorithm determines the features, with no influence of a priori knowledge of the sonar signals. The results of the two techniques, grouped by each of the twelve objects, are shown in Figure 2. In each of the twelve cases, the Bayesian network classifier performed as well as or better than the BCM-based classifier.

5. Conclusions and future work

While our results are preliminary, they would suggest that the addition of expert knowledge of the sonar signals at the least does not degrade classifier performance, and may improve it. It is likely that ultimately, the most effective means of implementing an automated classifier will combine unsupervised methods to extract information that is not readily discernible from the data with a means of utilizing a priori knowledge. This is particularly likely in the problem of mine detection, as there is ample body of expert knowledge to exploit. Ongoing work will refine the Bayesian network technique, and explore various hybrid approaches in order to develop an automated classifier that maximizes accuracy and efficiency.

References

- [1] J. Pearl. *Probabilistic Reasoning in Expert Systems: Networks of Plausible Inference*. San Mateo, CA, Morgan Kaufman, 1988.
- [2] R.A. Howard & J.E. Matheson. *Influence Diagrams, Principles and Applications of Decision Analysis*, Menlo Park, CA, Strategic Decisions Group, 1981.
- [3] D. Geiger, T. Verma & J. Pearl. *d-Separation: From Theorems to Algorithms, Uncertainty in Artificial Intelligence 5*, Amsterdam, North-Holland, 1990.
- [4] M.J. Larkin. Feature extraction in acoustic signals using the BCM learning rule, *Proceedings of the IEEE Asilomar Conference on Signals, Systems & Computers*, Pacific Grove, CA, pp. 889-893, November 1995.
- [5] E.L. Bienenstock, L.N. Cooper, & P.W. Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex, *Journal of Neuroscience*, 2:32-48, 1982.

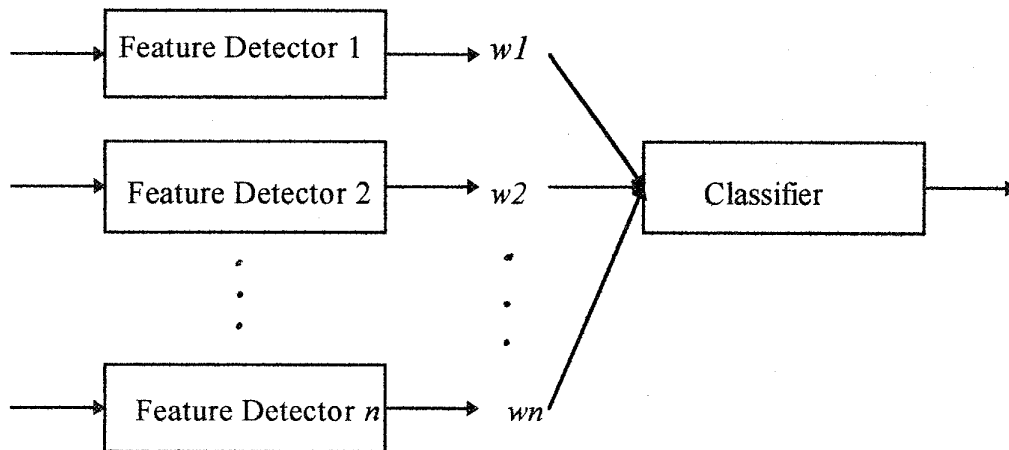


Figure 1. Schematic of Bayesian Network Classifier




Figure 2. Classification Results with Bayesian Network Classifier and BCM Classifier