

2013

Automated Detection of Vehicles with Machine Learning

Michael N. Johnstone
Edith Cowan University

Andrew Woodward
Edith Cowan University

Follow this and additional works at: <https://ro.ecu.edu.au/ism>



Part of the [Information Security Commons](#)

DOI: [10.4225/75/57b65924343cd](https://doi.org/10.4225/75/57b65924343cd)

11th Australian Information Security Management Conference, Edith Cowan University, Perth, Western Australia,
2nd-4th December, 2013

This Article is posted at Research Online.

<https://ro.ecu.edu.au/ism/161>

AUTOMATED DETECTION OF VEHICLES WITH MACHINE LEARNING

Michael N. Johnstone and Andrew Woodward
ECU Security Research Institute,
Edith Cowan University, Perth, Australia
a.woodward@ecu.edu.au; m.johnstone@ecu.edu.au

Abstract

Considering the significant volume of data generated by sensor systems and network hardware which is required to be analysed and interpreted by security analysts, the potential for human error is significant. This error can lead to consequent harm for some systems in the event of an adverse event not being detected. In this paper we compare two machine learning algorithms that can assist in supporting the security function effectively and present results that can be used to select the best algorithm for a specific domain. It is suggested that a naïve Bayesian classifier (NBC) and an artificial neural network (ANN) are most likely the best candidate algorithms for the proposed application. It was found that the NBC was faster and more accurate than the ANN for the given data set. Future research will look to repeat this process for cyber security specific applications, and also examine GPGPU optimisations to the machine learning algorithms..

Keywords

Security, Detection, Machine Learning, Optimisation

INTRODUCTION

Security professionals have at their disposal a range of devices that support the security function, for example, CCTV and X-ray scanners. One way in which computer science can assist in detecting security threats is through the application of machine learning in security detection devices and technology. Such application of machine learning to increase detection rates has been in use for over 50 years in the biomedical field (Ahmed, 2005). In particular, the use of artificial neural networks (ANN) in X-ray and computerised tomography (CT) scans for improved detection in the medical field has led to increased detection rates for cancer and other illnesses (Zhou and Jiang, 2002; O'Halloran et al., 2011; Ahmad and Khan, 2012). There is also evidence in the literature of the application of ANNs in the security field for increased detection of packages in baggage screening (Singh and Singh, 2004). Further reports describe the use of Bayesian models in the cyber security field to analyse images in emails to determine whether they are spam emails (Guzella and Caminhas, 2009).

The idea of analysing vehicle sounds is not new. Thomas and Wilkins (1972) noted that many factors contribute to the overall sound including engine noise, exhaust, wheels and air buffeting. Ding et al. (2004) proposed an adaptive threshold algorithm for real-time vehicle detection. Duarte and Hu (2004) used a k-nearest neighbor (KNN) approach to classification in a sensor network. Similar to Duarte and Hu, Malhotra et al. (2008) used a KNN approach with FFT data. Maciejewski et al. (1997) used neural networks for vehicle recognition, one approach we explore further in this paper.

Preliminary research in this area has identified candidate algorithms for use with classifying and identifying motor vehicles based on audio samples (Johnstone and Woodward, 2012). In this work we focus on real-time automated detection and analysis of sound-specifically motor vehicle noises. This has particular benefits for detection of criminal activity at remote locations as the system can be trained to detect classes of vehicle or even specific vehicles. This method has advantages over more conventional means of identifying and tracking vehicles as it does not rely on a device being attached to a vehicle nor does it rely on an occupant carrying a device that transmits a signal (such as a mobile phone). This paper tests two machine learning algorithms, namely artificial neural networks (ANN) and naïve Bayesian classifiers (NBC) using audio samples captured from motor vehicle engine sounds.

THE ANALYSIS OF SOUND SAMPLES

Capturing sound is a straightforward process, but analysing sound using a computer in place of a person is not as easy, and presents some not insurmountable difficulties. However, much depends on the domain of interest. For example, the requirements for speech recognition are different from music sampling. This is due to the way in which these sound types are produced, and whether the resultant wave forms are shown as represented in the

time, as opposed to frequency, domains (Figure 1). Whilst the time domain gives a good illustration of amplitudes and allows for comparison or contrast of a waveform over time, the frequency domain is more often used for analysis of audio signals because it highlights frequency peaks of interest that can facilitate comparison.

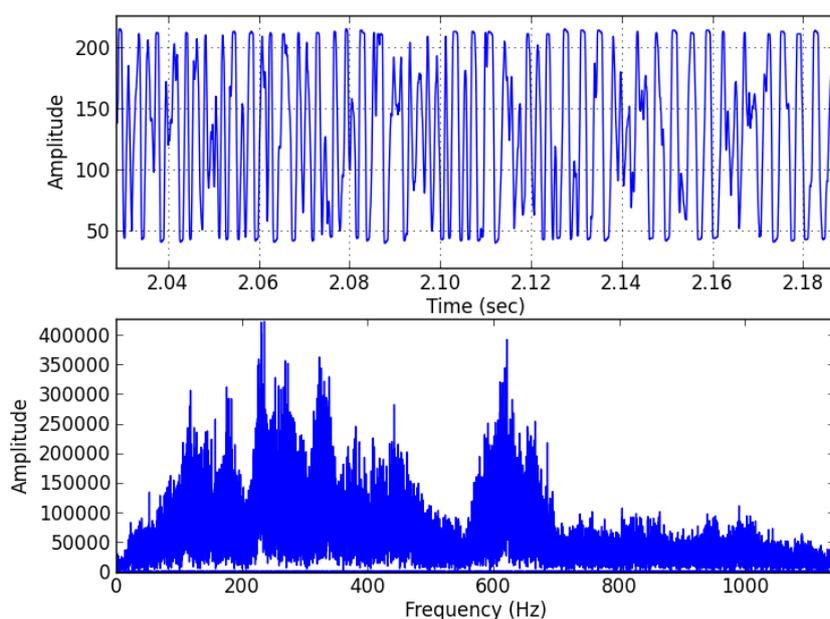


Figure 1. A Ford Mustang V8 engine sound wave sample represented in both time (a) and frequency (b) domains.

A major issue with the computational analysis of audio data is that such data are time-based. Sound is a wave transmitted through a medium such as air or water, and which consists of both amplitude and frequency varying over time. The wave form itself could be produced artificially by modulating one or more of the aspects of a wave to vary characteristics. This means that sounds which are of different lengths (times) are classified as different because their digital representation is different. This would not be a problem for a human operator, but can present difficulties for a machine learning algorithm.

A further issue is that of “aliasing” in digital signal processing. This issue can be avoided by sampling a signal at more than twice the frequency of the highest frequency of interest (i.e. the Nyquist frequency). Therefore, if the highest frequency of interest is 1000Hz, then as long as the sound is sampled at 2000Hz or more, then a frequency of 1000Hz can be analysed.

Having examined some of the issues involved in sound sampling and analysis, it is now appropriate to discuss algorithms that could be used to automatically classify sounds.

MACHINE LEARNING ALGORITHMS AS AUTOMATED CLASSIFIERS

Clearly, any algorithm that is used for the purpose intended here should be able to classify sounds and to learn new sounds based on some set of pre-specified criteria, much as a person would. A human can detect the difference between a truck and a car or a diesel or petrol engine easily. Asking a computer to do the same task presents some problems. Computers are swift at calculation, but not particularly good at drawing inferences from incomplete data (which of course is not a problem for an experienced security analyst). Search algorithms, particularly those that mimic human thought processes or behaviour, are likely to be of most use here. Specific types of search algorithm that have value in this context are naïve Bayesian classifiers and artificial neural networks (Graupe, 2007), and there is evidence in the literature of such algorithms being applied to audio (Krishnamoorthy and Kumar, 2011).

A naïve Bayesian classifier (NBC) assumes that the determining attributes of a class are independent of one another. This means that using such an approach is swift and requires little training. The assumption may not be true in practice (and often isn't—the attributes that determine an object are frequently dependent on one another) but can be accurate despite this simplification.

An artificial neural network (ANN) is a structure that mimics the neurons inside a brain. At its simplest, a single network node may have multiple inputs and the node only fires the output provided that the sum of the inputs reaches a predefined value. Recall that we are attempting to use machine learning to provide a substitute

for a human operator for some aspects of the security function. Given that aim, neural networks appear to show promise as an effective and efficient machine learning algorithm (Rumelhart and McClelland, 1986).

Essentially the problem to be solved is classification, therefore algorithms for both ANNs and NBCs can provide legitimate solutions as both are used for classification purposes. A key assumption of ANNs is that the input variables are linearly separable, although the input can be non-parametric. We adhere to this assumption by using a frequency band approach. Bayesian methods are not nearly so sensitive to overlapping input as ANNs, but the former do assume a Gaussian distribution for the data. As such, a Gaussian based NBC approach was used for this research.

EXPERIMENTAL METHOD

We used a Python program to generate the Fast Fourier Transforms of the initial .wav files for three motor vehicles. Clearly, the waveforms are too complex to be easily interpreted (Figure 1a) thus the frequency spectrum (Figure 1b, bottom) for each vehicle was analysed for significant (large-amplitude) peaks. The peaks were then assigned to frequency bands between 0-1000Hz. These frequency profiles be-came the input variables to the learning systems.

In order to provide training data, each frequency profile was randomly varied by $\pm 1\%$ (or 10Hz) to account for Doppler shifting as a vehicle approached a sound recorder, and then receded. One hundred cases were generated for each vehicle. A test data set, comprised of a further 100 cases for each vehicle was also created. The frequency profile of the test set was randomly varied by $\pm 10\%$ (or 100Hz) in order to provide a challenging data set where some frequencies would fall into the wrong band and thus the output would likely be misclassified. Both data sets were scaled as per advice provided in Haykin (2009) prior to analysis. Any negative frequencies resulting from the random variates were set to zero.

The test system was a 1.6GHz Core 2 Duo system with 2 GB of RAM, running Python 2.6. Whilst this system would appear to be slow in comparison with most cur-rent-generation PCs and certainly with any reasonably-optional GPGPU-based sys-tem, the target processor has a small footprint in terms of speed and processing capacity, therefore using a (relatively) low-power machine as an experimental test bed was justified.

RESULTS AND DISCUSSION

Testing an artificial neural network

The first learning system to be tested was an artificial neural network (ANN). The data described in the previous section were input to an ANN with a single hidden layer.

The performance of the ANN was quite good, even with the wide-variation test set. With this data set, the success rate was 83.7% with a hidden layer consisting of three nodes. An expanded hidden layer of 8 nodes boosted the success rate to 88.7%. Given that in the test data set the frequencies were deliberately set to contribute to mis-classification, this is a very good success rate. Figure 2 shows the (well-behaved) error plot for a typical dataset used in these experiments. As a further test, another data set with more well-behaved data (1% random variation for both training and test sets) was also analysed, with a resultant success rate of 100%. Therefore it can be concluded that the ANN algorithm is effective in solving this classification problem. A discussion of efficiency follows in the next section.

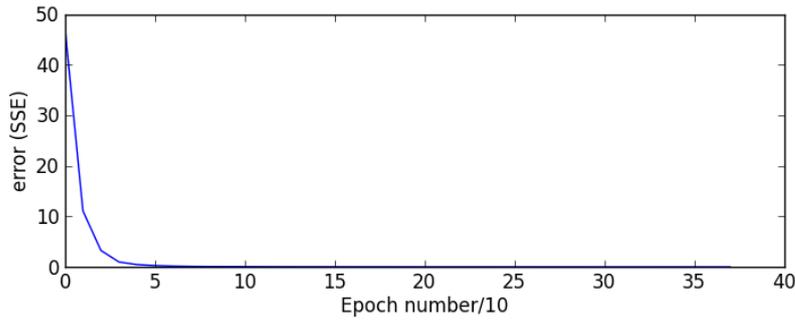


Figure 2. Error Variation in ANN Training.

Factors affecting the Performance of the ANN

We used a conventional backpropagation algorithm, modified from Blais and Mertz (2001). Several factors influence the performance of the algorithm, including the algorithm mistaking a local minimum for the global minimum and the number of neurons in the hidden layer.

Figure 3 shows a section through an error plane. The error surface is multivariate, but it is instructive to view it in two dimensions nonetheless. The algorithm works by looking for the steepest gradient and traversing that pathway until the minimum is reached. As figure 3 shows, if the initial point is near a local minimum (z_1 or z_2), the algorithm will gravitate towards that local minimum and not towards the global minimum (z_{min}). Figure 4 shows, however, that the algorithm can move the system from a local minimum to the global minimum. The algorithm has problems initially but then settles after 400 iterations. This leads to a design decision regarding whether to modify the algorithm to discard a (potentially optimal) minimum or to consider satisficing i.e. that any minimum is sufficient, even if it is not the global (most desirable) minimum. Recall that the error surface is multidimensional, so any attempt to search the entire surface for the global minimum will take considerable time. The approach taken in this research is to search for the nearest minimum.

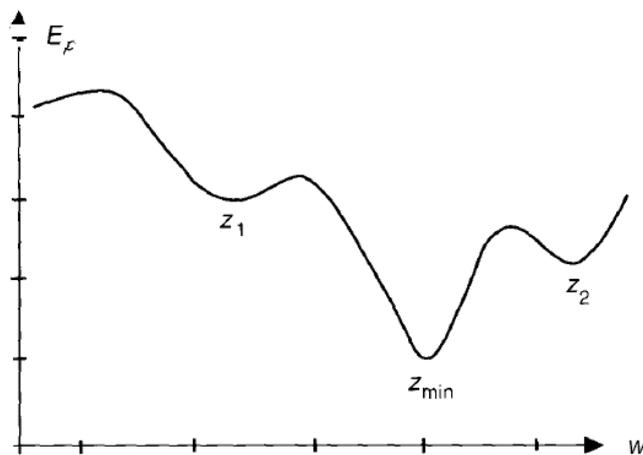


Figure 3. Section through an Error Hyperplane (adapted from Freeman and Skapura, 1991).

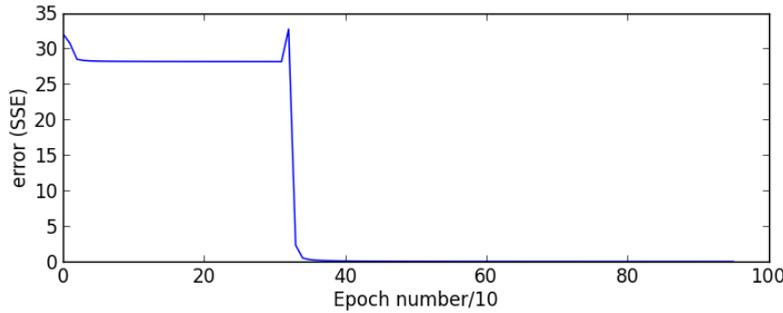


Figure 4. Error Variation in ANN Training with Local Minimum Problem.

We found that adding more neurons to the hidden layer increased the training time of the system considerably and that the relationship between the number of hidden nodes and the training time to be linear (at least over the range of interest). What was important here was that whilst this increased training time (an undesirable outcome), it also allowed the system to come to an equilibrium with respect to the error more swiftly (a desirable outcome). Thus, there was a balance between the training time, the number of neurons in the hidden layer and the speed of convergence to an error minimum.

The training time was 26.3 seconds for a dataset of 600 training and test records, with 1000 iterations. In contrast, the testing time was approximately 0.03 seconds. By adjusting the number of neurons in the hidden layer and the number of iterations with respect to the error term, it was possible to reduce the training time to 7.1 seconds with no loss of accuracy (i.e. a constant error term and success rate). The reduction in training time achieved by tuning the system shows that the ANN algorithm can be efficient.

Testing a naïve Bayesian classifier

There are three main event models which can be used when applying a NBC to a dataset for the purposes of classification, and these are dependent upon the distribution of the data. For the purposes of classifying discrete datasets, multinomial and Bernoulli distributions are often employed. For continuous data, which is the target of this study, a Gaussian distribution is most appropriate, and was used in this study.

A training and testing regime consisting of 300 training and 300 testing samples, with 1000 test iterations, was used. This resulted in a training time of 0.107 seconds and 92.3% accuracy rate. A comparison of results for the same data set between the NBC and ANN algorithms is presented in Table 1.

Table 2. Summary of Algorithm Performance.

Algorithm	Training time (sec)	Test time (sec)	Percent classified correct
ANN	7.102	0.031	88.7%
NBC	0.107	0.005	92.3%

CONCLUSIONS AND FURTHER WORK

This study sought to explore whether machine learning algorithms could assist or supplant human analysts in detecting specific motor vehicle sounds, thus providing a more effective security function in terms of potentially discovering criminal activity. The complex nature of sound sampling was revealed and the results of a series of experiments that tested implementations of two algorithms were articulated and discussed.

Specifically, this study tested an artificial neural network and a naïve Bayesian classifier. It was found that, for this application, the naïve Bayesian classifier outperformed the artificial neural network. Also, both algorithms can reliably detect different types of vehicles much faster than a human can (in the order of hundredths or thousandths of a second), thus suggesting that these algorithms can replace a human operator successfully.

A limitation of this work is that it used data from only three vehicles, therefore it should be considered proof-of-concept. Further work would involve extending the number and type of vehicles as well as examining one specific make and model of vehicle.

In addition, a number of other directions present themselves for this research. Firstly, in addition to simply examining the use of assisted machine learning for application to classifying audio signals, the scope will be broadened to include examples where it is used in cyber security. For example, email filtering and intrusion detection systems. A further research direction and an expansion of both the original research discussed here and of the examination of algorithms used in cyber security applications, is to look at the efficacy of each algorithm using a general purpose graphics processing unit (GPGPU). Such GPGPUs are currently being used extensively in testing encryption strength, as they lend themselves to FPU intensive calculating tasks. Future re-search will conduct a comparison of CPU vs. GPGPU for each algorithm type and application, the aim being to produce a framework which allows for the most appropriate and best performing algorithm to be selected for a given cyber security application. There is also the possibility of extending this to compare the same characteristics but with the two variables being clustered CPUs vs. clustered GPGPUs.

REFERENCES

- Ahmad, A. M. and Khan, G.M. (2012) Breast cancer detection using cartesian genetic programming evolved artificial neural networks. Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference. Philadelphia, Pennsylvania, USA, ACM: 1031-1038.
- Ahmed, F. E. (2005). "Artificial neural networks for diagnosis and survival prediction in colon cancer." *Molecular Cancer* 4(29): 12.
- Blais, A. and Mertz, D. (2001). An introduction to neural networks: Pattern learning with back-propagation. http://gnosis.cx/publish/programming/neural_networks.htm .
- Ding, J., Cheung, S.-Y., Tan, C-W and Varaiya, P. (2004). Signal processing of sensor node data for vehicle detection. Proceedings of the 7th International IEEE Conference on Intelligent Transportation Systems, 2004. pp70-75.
- Duarte, M.F. and Hu, Y.H. (2004). Vehicle Classification in Distributed Sensor Networks. *Journal of Parallel and Distributed Computing* (64), pp826-38.
- Freeman, J.A. and Skapura, D.M (1991). *Neural Networks: Algorithms, Applications, and Programming Techniques*. : Addison-Wesley, Reading, MA.
- Graupe, D. (2007). *Principles of Artificial Neural Networks*. 2nd Ed. Singapore: World Scientific.
- Guzella, T. S. and Caminhas, W.M. (2009). "A review of machine learning approaches to Spam filtering." *Expert Systems with Applications* 36(7): 10206-10222.
- Haykin, S. (2009). *Neural Networks and Learning Machines*. 3rd. ed., Pearson, Upper Saddle River, NJ.
- Johnstone, M.N., and Woodward, A. (2012) 'Towards effective algorithms for intelligent defense systems'. In: Yang Xiang, Javier Lopez, C.-C. Jay Kuo, and Wanlei Zhou, (eds.) *CSS 2012. LNCS*. vol. 7672, pp. 498-508. Springer, Heidelberg.
- Krishnamoorthy, P. and Kumar, S. (2011). Hierarchical audio content classification system using an optimal feature selection algorithm. *Multimed. Tools Appl.* 54:415–444.
- Maciejewski, H., Mazurkiewicz, J., Skowron, K. and Walkowiak, T. (1997). Neural Networks for Vehicle Recognition. Proceedings of the 6th International Conference on Microelectronics for Neural Networks, Evolutionary and Fuzzy Systems, pp5-9.
- Malhotra, B., Nikolaidis, I. and Harms, J. (2008). Distributed classification of acoustic targets in wireless audio-sensor networks. *Computer Networks*. 52(13), pp2582-2593.
- O'Halloran, M., Mcginley, B., Conceição, R.C., Morgan, F., Jones, E. and Glavin, M. (2011) "Spiking neural networks for breast cancer classification in a dielectrically heterogeneous breast." *Progress In Electromagnetics Research* 113: 16. pp. 413-428.
- Rumelhart, D.E. and McClelland, J.L. (1986) *Parallel distributed processing: Exploration in the microstructure of cognition*. Cambridge, MA: MIT Press.
- Singh, M. and Singh, S. (2004). "A knowledge-based framework for image enhancement in aviation security." *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on* 34(6): 2354-2365.

- Thomas, D.W. and Wilkins, B.R. (1972). The analysis of vehicle sounds for recognition. *Pattern Recognition* (4), pp379-89.
- Zhou, Z.-H. and Jiang, Y. (2002). "Lung cancer cell identification based on artificial neural network ensembles." *Artificial Intelligence in Medicine* 24(1): 25-36.