Optimal Detection and Resource Utilisation for the Detection of Anomalous Vehicle Behaviour at Checkpoints

J.D. Rogers and M. Bernhardt
Waterfall Solutions, Parklands,
Guildford, Surrey GU2 9JX.

Abstract

The detection of anomalous vehicle behaviour is of crucial importance to military security, as it can be indicative of potentially hostile actions. Consequently, the development of automated methods for aiding the detection of anomalous behaviour would be hugely beneficial. Although the use of passive Electro-Optic (EO) sensors for directly detecting objects that are hidden within vehicles may not prove to be useful, the problem can potentially be solved through an analysis of vehicle behaviour. The method implemented within this work uses EO sensors to detect and track vehicles, and exploits recent advances in conformal prediction theory to provide a calibrated confidence framework and an optimal detection rate.

Keywords: Tracking, Anomaly Detection

Introduction

In current operations the threat from attack by adversaries using devices hidden within vehicles is substantial. Due to the nature of these devices, their detection is an extremely complicated problem and is a topic of much current research. The task is approached here from a behavioural analysis viewpoint. The behaviour of vehicles that are carrying anomalous devices could signify crucial information concerning their intent and provide hugely beneficial early warnings.

This project [1] was conducted by Waterfall Solutions Limited (WS) on behalf of the EMRS DTC and had the aim of enhancing the detection of anomalous vehicles at checkpoints. The project combined algorithms for vehicle detection and tracking, as well as anomaly detection, with recent developments in the theory of conformal prediction. A primary consideration in this research was optimising detection performance, as each potential detection event is likely to lead to the need to search a vehicle in detail, which is a potentially risky and resource-intensive action. A high number of false-positive alarms will cause unwanted tail-backs at checkpoints, lead to operator mistrust of the system, and result in poor use of available resources.

The method developed within this work employs a novel feature tracking method to convert vehicle tracks into a more useful representation, to which anomaly detection methods can be applied. The vehicle tracks were obtained through extended object detection methods which were also investigated within this work. The anomaly detection technique was implemented within the conformal prediction framework to provide valid confidence estimates in the predictions and result in an optimal detection rate.

Extended Object Tracking

The first component of this work concerned the detection and tracking of vehicles within realistic environments. A number of different methods exist for tracking extended objects. One of the most basic methods for detecting movement within video imagery is temporal differencing, where the possible detections are defined as those portions of the image which differ significantly to the previous image. The difference between frames I_t and I_{t-1} for the pixel at position x and y is given by:

$$\delta_t(x, y) = abs[I_t(x, y) - I_{t-1}(x, y)]$$

Due to the homogeneous nature of some objects, this method may tend to only detect the edges of moving objects because these are the only regions of the image that are substantially changing. However, this method can be improved by averaging these differences over a number of frames to alleviate the problem. This simplistic approach can prove to be useful for many situations. However, one of the drawbacks of this detection method is that if the vehicle is stationary for a period of time, the detection method will fail and the track is likely to be lost. Another problem with this method is that it can be misled by natural fluctuations in the imagery. These fluctuations are typified by the movements of trees and bushes in the scene.

Figure 1 illustrates the result of applying the temporal differencing approach to a sequence of vehicle imagery. It can be seen that the detected clusters for the vehicles are well defined, although there are many false alarms associated with the bushes in the left of the image.

Another technique for detecting motion within imagery is to calculate the optical flow in the scene. This class of technique identifies movement of objects within imagery and therefore detects moving regions rather than changing edges. The magnitude of the movement can then be used as a measure of detection, which provides a more robust technique for extended object detection because it is less likely to be misled by pixel changes caused by natural fluctuations in the imagery. The flow field can also be used to provide additional information about the vehicle motion. Similar to the temporal differencing method, this technique is based on temporal changes and will lose the track if the vehicle stops for a period of time.

Figure 1 – Result of applying temporal differencing to vehicle sequence. Blue components indicate output of the detection, orange components are the accepted clusters (after morphology) and the pink components are the overlap of the two.

Figure 2 shows the result of applying the optical flow method to the same sequence as before. The figure indicates that the optical flow technique is more robust than the temporal differencing approach to the small movement in the trees and bushes. However, the clusters that are obtained for the vehicles are not as well defined because of the nature of the flow field, which is constructed using a combination of a windowing technique and a multi-resolution approach.

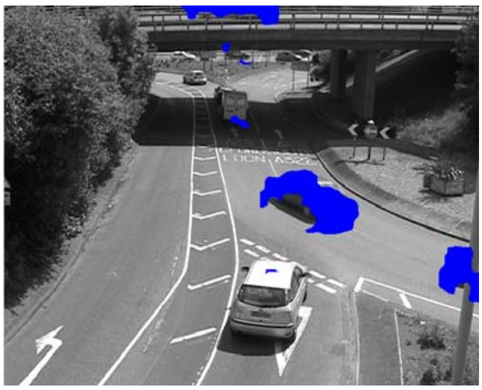

Figure 2 - Result of applying optical flow to vehicle sequence. Blue components indicate output of the detection.

One method that is capable of continuing to detect vehicles when they are stationary is background subtraction [2]. This technique builds a model of the background and then highlights any part of the image that differs greatly from this model. This technique can build natural fluctuations in the imagery into the background model and is therefore more robust to such changes. The background values of each pixel can be modelled by a uni-modal Gaussian distribution, where a pixel is defined as foreground if it differs from the mean by more than a fixed number of standard deviations.

Figure 3 illustrates the result of applying the background modelling technique to the vehicle sequence. It can be seen that background subtraction is the only method that continues to track the stationary vehicle.

Figure 3 – Result of applying background subtraction to vehicle sequence. Pink components are the output of the detection.

One technique that was implemented in this project combined the benefits of the above three methods into a single detection technique. The algorithm calculated the flow magnitude of the image, and clustered regions of high flow to form a detection mask. This mask was then used to protect the background model for foreground detection. The background pixels were only updated if they were not considered to be foreground or part of a high flow region. The flow field was also used to establish new tracks. If a region of the image is detected as foreground and cannot be associated to an established track, the flow of that cluster is determined such that the cluster is rejected if it is not a region of high flow. This method will therefore only initiate a track if it is an object with a large degree of motion.

Anomaly Detection

The process of anomaly detection concerns the identification of examples which appear to contradict the general trend of the data. It is a problem which can be approached through the estimation of the underlying data distribution and the comparison of new examples to this distribution.

For multi-variate data, a measure of how atypical an example is in regard to a distribution can be computed from the Mahalanobis distance. However, this measure assumes that the data constitutes a uni-modal distribution and, therefore, is not as effective if the data is multi-modal. In this work, the anomaly detection algorithms are applied to vehicle tracks. The tracks to be considered may belong to one of a number of distributions, and the anomaly detection method should learn these different distributions and define an anomaly as an example that is unlikely to belong to any of the distributions that it has previously seen. Clustering methods are suitable for this purpose as they group nearby examples to form separate clusters which can then be examined. The method implemented within this work employs a tree structure to cluster the data, and then compares new examples to this model.

Most machine learning algorithms output ‘bare’ predictions: that is, predictions without any associated confidence measure. A methodology is described in [2] that allows any prediction algorithm to be modified in order to output predictions that are also equipped with a confidence measure. Alternatively, a required degree of confidence can be stipulated as an input parameter and the algorithm will output a ‘tolerance region’ that will contain the true answer with the required degree of confidence. These confidence measures are proved to be valid under a fairly simple condition: that the input data are independent samples from some unknown distribution.

All previously published work in Conformal Prediction theory has concentrated on supervised learning where there are a number of different classes available. With some modifications, the theory can also be applied to the case of a single class and the question of finding anomalies within data. Instead of a classifier, one now has an anomaly detector that evaluates how similar new samples are to previously seen samples. The output of this anomaly detector can then be directly interpreted as a non-conformity score. The theory then allows the assertion that the errors at a given confidence level ε will be independent, identically distributed, and occur with a probability of precisely 1- ε. This is not at all obvious, since the anomaly detector itself may be updated 'on the fly', and continue to learn from new samples. The non-conformity score employed within this work is based on the structure of the clustering tree.

Bayesian methods are an alternative class of anomaly detection algorithm, which assumes that the prior distribution of the data is known. This will actually be optimal if the prior distribution is correct. In practice, this is almost never the case since prior distribution models are chosen at least partially on the basis of computational expediency. It has been shown [4] that Bayesian algorithms with marginally incorrect priors perform less well than conformal prediction which makes no prior assumptions.

Alternatively, other methods employ a training set of data, from which a static threshold can be calculated. This technique may require a large amount of data to produce reasonable threshold estimates. It is also essential that this training set is acquired under exactly the same conditions as those where the anomaly detector will eventually be used. In the case of conformal anomaly detection, this learning is performed on-line and so each system will adapt to its local conditions rather than using a large 'average conditions' data set.

Another typical method for anomaly detection will detect a fixed fraction of materials within a given time window. This is wasteful of resources since there will be times when very few vehicles display anomalous behaviour, yet this approach would still choose a constant fraction to be searched. Similarly, if there was an unusually high concentration of anomalies this method might miss some. By allowing the algorithm to adapt in an on-line manner in the conformal approach, it can maintain an overall constant alarm rate *on average.* This allows for periods when there are few anomalies detected (implying that few vehicles get searched) and periods when it will correctly recommend the searching of many vehicles due to the relatively high concentration of anomalies.

Results

Using the vehicle tracking methods, a set of eleven vehicle tracks were obtained. These can be separated into two groups: those which travel in the left lane and follow the road to the left; and those which travel in the right lane and follow the road to the right.

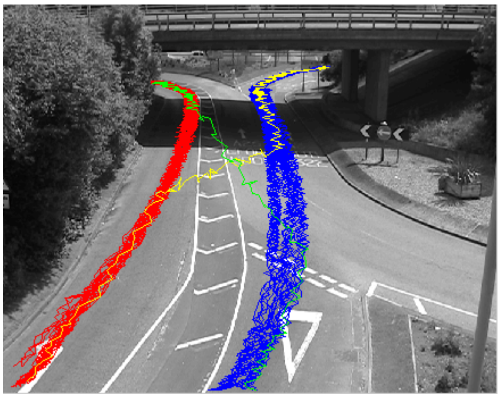

Figure 4 - Plot illustrating 50 standard tracks and 2 anomalies. The standard tracks are separated into the two classes indicated by red and blue. The two anomalies are shown in green and yellow.

By adding random noise to the tracks a larger set of representative tracks were created, as shown in Figure 4. Anomalous tracks can also be created through combining one track from each set and interpolating between them. Two of these anomalies have been created and added to the figure.

This work employed WS' novel track feature representation to convert these tracks into a more useful space. This feature representation yields the plot of Figure 5. It can be seen from the figure that projecting the tracks onto features 1 and 2 is sufficient to identify the two anomalies.

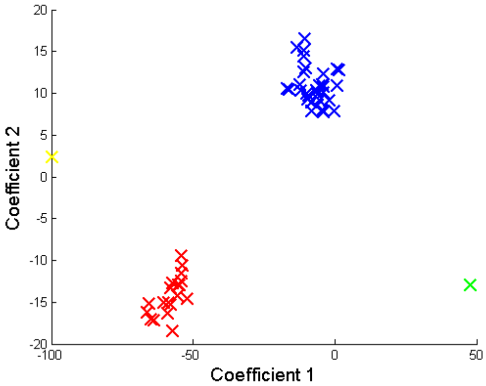

Figure 5 - Feature representation of modified real tracks using coefficients 1 and 2. The two standard track classes are shown in red and blue, and the two anomalies are shown in yellow and green.

In order to generate a large data set, 1000 of these tracks were created including 10 anomalies. The tracks were then converted into the feature representation and shown sequentially to the anomaly detector.

As each new example was observed, the anomaly detector applied the clustering approach to all of the data it had seen up until that time in order to extract conformity values for each of the examples. The confidence level was set at 0.05 so that if the conformity measure fell into the 0.05 least conformal proportion of the data it could be declared as an anomaly. Figure 6 shows the plot of the cumulative detection rate against the number of incoming vehicles. It can be seen from the figure that the probability of detection remains approximately constant as the algorithm proceeds.

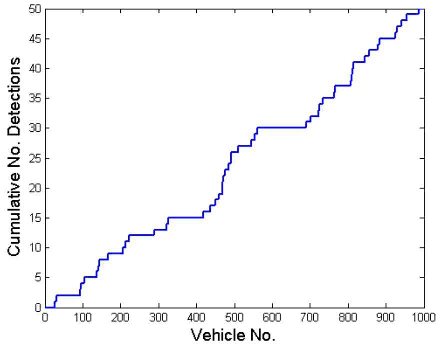

Figure 6 - Graph showing cumulative detection rate against time.

Figure 7 illustrates the histogram showing the number of detections within a window of 50 vehicles. With the confidence level set at 0.05, the expectation for the number of detections in the time window of 50 vehicles should be 2.5 on average. The number of detections within a window should be independent of the time since the last detection and if this is true, then the histogram should have a Poisson distribution with a rate of 2.5. The expected values of the corresponding Poisson distribution are also plotted on the figure and can be seen to closely match the observed values. This clearly shows that the detections occur independently and at the correct average rate.

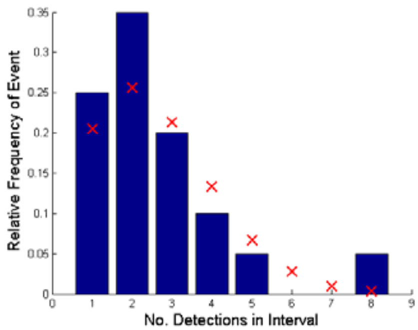

Figure 7 - Histogram of number of detected anomalies and the predicted quantities from the corresponding Poisson distribution, which are shown with red crosses.

The confidence level was set at 0.05 which resulted in 50 detections. The 10 synthetic anomalies that were added to the data were correctly within this set of 50.

Discussion

The work that has been reported here has established the successful detection of anomalous vehicle behaviour. As a precursor to behavioural analysis, the vehicles must first be detected and tracked, and so a number of approaches to vehicle detection and tracking were tested on appropriate recorded imagery. The results of these different detection and tracking techniques (taken from WS' Versatile Image Processing Application toolbox) have been presented and discussed.

Real image data was processed using extended object detection methods and the subsequent tracking methods were able to extract a large number of tracks. Since the real data did not contain any unusual behaviour, these tracks were modified to create a large data set containing several anomalies. The synthesised tracks were converted into the track feature space and delivered sequentially to the anomaly detection system. The anomaly detector was based on a clustering tree to enable the technique to work on multi-modal data. This capability was demonstrated on the track data as this data consisted of two sets of typical tracks. All of the synthetic anomalies were detected by the algorithm. This method updates the detection algorithm in an on-line manner and adapts the technique accordingly. This is done optimally as anomalies were predicted at the correct rate and independently of the time since the last detection.

Although this work has demonstrated successful vehicle tracking, it is recommended that more research is done to develop a more robust background model for vehicle detection. This would then provide reliable track data in real scenarios. It would also be interesting to investigate the combination of this passive sensor method with direct detection approaches, where this technique could be used to guide the direct techniques to examine vehicles exhibiting unusual behaviour.

References

1. Waterfall Solutions Ltd, "Optimal Detection and Resource Utilisation for IED Discovery at Checkpoints: Final Report", WS/08/EMRS/IED/FR/001, 31st July 2008.
2. M. Piccardi, "Background Subtraction Techniques: A Review", *Proc. IEEE International Conference on Systems, Man and Cybernetics,* 2004.
3. V. Vovk, A. Gammerman and G. Shafer, "Algorithmic Learning in a Random World", Springer, March 2005.
4. Technical Report Royal Holloway University of London, "Conformal Anomaly Detection", July 2008.

Acknowledgements

The work reported in this paper was funded by the Electro-magnetic Remote Sensing (EMRS) Defence Technology Centre, established by the UK Ministry of Defence and run by a consortium of SELEX Galileo, Thales UK and Roke Manor Research