University Defence Research Collaboration in Signal Processing

Edinburgh Consortium White Paper

Developing Automated Anomaly Detection for Wide Area Surveillance (WAS)

Introduction

Surveillance data is commonly used for situation awareness for both civilian security and defence applications. One of the key benefits of visual surveillance is that areas can be monitored remotely, and this is one property that is providing new surveillance opportunities - particularly on the battlefield. With the increasing prevalence of both manned and unmanned aircraft equipped with high resolution sensors, the ability to surveil large areas brings numerous technical challenges.

Of particular interest to us is how we can process WAS signals in such a way as to identify the salient parts of the data both online and in real-time. For example, existing wide area motion imagery systems are capable of capturing 210 Megapixels/second and may cover areas from 6 – 50 KM2 [1]. Automated algorithms are desperately needed to reduce analyst information overload. Furthermore, the provision of automated processing paves the way for a pro-active, rather than reactive capability.

Techniques for automated surveillance are often based on anomaly detection algorithms. The key insight behind anomaly detection is that for most surveillance tasks examples of ‘uninteresting’ (normal) behaviour are abundant, while examples of actual events of interest are scarce and hard to define. Anomaly detection tackles these problems by inverting the problem – that is – by attempting to model normal behaviour, abnormal (anomalous) behaviour can be identified by its poor fit. Of course, the anomalies detected are only as reliable as the underlying model of normal behaviour, and thus a detected anomaly is only an indicator that the underlying model cannot explain the behaviour observed. Whether that behaviour is genuinely ‘anomalous’ or just infrequent must still be determined, and thus completely automated surveillance is challenging to achieve. Nevertheless, the ability to filter surveillance data by saliency is still an important capability, particularly for WAS, where the number of targets concurrently observed could be in the thousands. Any algorithm that can reduce in the number of candidates offers clear potential for reducing information overload.

Despite wide interest in developing automated surveillance techniques, very little work has addressed the challenges that must be overcome before WAS based anomaly detection can be achieved. Not least of these is the sheer volume of data that must be processed and represented in an efficient way, with many techniques focusing on small scenes with few moving targets. Online and adaptive algorithms for learning common target motion are also scarce within the literature, and the ability to detect anomalies from partial target trajectories – as they evolve – is rarely considered.

EPSRC
Engineering and Physical Science Research Council

[dstl]

UDRC

Method

As part of Phase 2 of the University Defence Research Collaboration, Heriot-Watt University is considering many of these problems within the context of wide area motion imagery (WAMI) surveillance while maintaining applicability to broader surveillance problems. A key principle of our approach is that normal behaviour can rarely be modelled globally, and thus both spatial and temporal context are of high importance. At the spatial level trajectory clustering techniques inspired by [2] and [3] are being developed to model where targets normally move. The vast quantity of observed data is 'wrapped up' into more compact distributions using chains of Gaussian distributions representing unique trajectory patterns. Not only can target behaviour be matched to existing clusters in real-time, but on-line learning is also performed allowing the underlying model to be refined and adapted in response to changing scene behaviour.

Many of the environments we would wish to monitor are too complex to be modelled using spatial context alone. By modelling temporal context we not only model where target activity normally occurs, but crucially, when. Using kernel density estimation techniques adapted from [4] temporal activity distributions can be learnt for each spatial cluster in a highly scalable way suitable for an on-line, persistent surveillance platform.

Combining these techniques is likely to provide a real-time and on-line anomaly detection algorithm that is capable of learning from - and detecting anomalies within – large streaming datasets such as WAMI. Moreover, this work will provide the first algorithm able to learn spatio-temporal motion patterns over large areas, making it suitable not only for anomaly detection of land, sea and air targets, but also capable of mapping activity in unknown regions. What's more, adopting the view that target trajectories are nothing more than 2-dimensional signals highlights the broader exploitation potential of our algorithm to other multi-dimensional signals.

References:

[1] A. Basharat et al., "Real-time multi-target tracking at 210 megapixels/second in wide area motion imagery". Applications of Computer Vision, IEEE Winter Conference on. 2014

[2] W. Hu et al., "A system for learning statistical motion patterns". Pattern Analysis and Machine Intelligence, IEEE Transactons on. 28(9), 1450-1464. 2006.

[3] C. Piciarelli and G. Foresti. "On-line trajectory clustering for anomalous events detection". Pattern Recognition Letters. 27(15), 1835-1842. 2006.

[4] M. Kristan et al. "Online discriminative kernel density estimator with Gaussian kernels". Cybernetics, IEEE Transactions on. 44(3), 255-265. 2014