



Anomaly detection in vessel tracks using Bayesian networks

Steven Mascaro^{a,*}, Ann Nicholson^b, Kevin Korb^b

^a Bayesian Intelligence Pty Ltd., 2/21 The Parade, Clarinda, Victoria 3169, Australia

^b Clayton School of IT, Monash University, Wellington Road, Clayton, Victoria 3800, Australia



ARTICLE INFO

Article history:

Available online 2 April 2013

Keywords:

Machine learning
Bayesian networks
Models of normality
Anomaly detection
AIS
Maritime data

ABSTRACT

In recent years electronic tracking has provided voluminous data on vessel movements, leading researchers to try various data mining techniques to find patterns and, especially, deviations from patterns, i.e., for anomaly detection. Here we describe anomaly detection with data mined Bayesian Networks, learning them from real world Automated Identification System (AIS) data, and from supplementary data, producing both dynamic and static Bayesian network models. We find that the learned networks are quite easy to examine and verify despite incorporating a large number of variables. We also demonstrate that combining dynamic and static modelling approaches improves the coverage of the overall model and thereby anomaly detection performance.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

A wealth of information on vessel movements has become available through the use of the Automated Identification System (AIS), with much of it even filtering through to the public via the Internet. Surveillance authorities are interested in using this data to uncover threats to security, illegal trafficking or other risks. Whereas previously surveillance has suffered from a lack of data, electronic tracking has transformed the problem into one of overabundance, leading to a need for automated analysis.

The main goal of vessel behaviour analysis is to identify anomalies. As noted by Riveiro and Falkman [1], anomalies are detected either by using signature-based approaches or, as we do here, by developing a model representing normal behaviour, with anomalous behaviour being then identified by the extent of a vessel's deviation from normality. A common approach to creating normal models is to cluster the data around a set of points in a multi-dimensional feature space, with features such as longitude and latitude, speed and course [2]. Tracks that are within or near one of these clusters are considered normal, while the remainder are flagged as potential anomalies. Researchers use many different machine learning techniques to generate normality models from vessel movement data (typically AIS data), including the learning of Gaussian mixture models [2], support vector machines [3] and neural networks [4]. A disadvantage of these approaches is that they do not provide a transparent model that a human user, such as a surveillance officer, can understand, interact with and explore.

Here, we explore the use of Bayesian Networks (BNs) [5,6] for analysing vessel behaviour and detecting anomalies. While BNs have been widely applied for surveillance and anomaly detection (e.g., [7–10]), to date there have been only a few preliminary applications of BNs to maritime anomaly detection. As noted by Johansson and Falkman [11], however, BNs potentially have two substantial advantages in this domain over other types of models: (1) BN models are easily understood by people who are not BN specialists (which may include surveillance operators or other domain experts) and (2) they allow

* Corresponding author.

E-mail address: steven.mascaro@bayesian-intelligence.com (S. Mascaro).

URL: <http://www.bayesian-intelligence.com> (S. Mascaro).

for the straightforward incorporation of expert knowledge. They can also represent causal relations directly and, in that case, have the advantage of being more easily verified and validated, as we show in Section 3. We begin with a brief look at some of earlier approaches to anomaly detection.

1.1. Other approaches to anomaly detection

Support Vector Machines (SVMs) partition the multidimensional feature space, producing strict boundaries between clusters. In their simplest forms, SVMs suffer from a number of problems that have limited their use in vessel anomaly detection, including a lack of partial assignment, a restriction to binary classes, high computational complexity and difficulties in summarizing and communicating the learned models. Li et al. [3], however, make use of SVMs to perform an interesting analysis of vessel behaviour at a higher level of abstraction than that of the time series. Li et al. extract higher level movement features from the track (such as turning left or looping) and then cluster these further into what they call “movement motifs”. They show that an SVM trained on the movement motif abstractions can correctly classify a significantly higher percentage of their test data in some cases than an SVM trained on lower level features alone.

One commonly used model is the neural network [4,12], which consists of a network of processing nodes, input/output connections between nodes and weights attached to the connections. For anomaly detection neural networks are typically used to map an input vector of reals to an output in the form of a classification. When used in this way, a neural network partitions the feature space much like an SVM. Unfortunately, data mined neural networks of any moderate degree of complexity are almost completely opaque to human understanding, whereas their interpretation by surveillance operators is one of their primary purposes [13].

Gaussian Mixture Models (GMMs) have proven a popular choice for representing normality models of vessel behaviour [14, 2,15]. As its name implies, a GMM is a combination of multi-variate Gaussian distributions. These distributions aim to summarize how the training data cluster and spread in the multi-dimensional space. Kernel Density Estimators (KDEs) are a generalisation of GMMs, using a sum of (typically Gaussian) distributions for each point, they allow for more flexibility than GMMs in the way clusters are described. Unfortunately, both GMM and KDE models can be difficult for non-experts to understand. Laxhammar et al. [15] trained both GMMs and KDEs on AIS data and evaluated anomaly detection performance by stochastically generating anomalous tracks, and then measuring how many steps it took for each method to flag the track as anomalous. They found little extra value in using KDE methods over GMMs.

Das and Schneider [16,17] identify anomalous cases by finding unexpected dependencies between sets of attributes. They compare their approach to one using BNs and find their approach does better. We are skeptical inasmuch as their approach could very readily be adopted using Bayesian networks, while Bayesian network models also allow the identification of a large further class of anomalies not reflected simply by direct dependencies, for example, those represented only by conditional dependencies, which Das and his collaborators ignore.

1.2. BN-based approaches to anomaly detection

Given that anomalies just are events that are highly improbable under ordinary circumstances, Bayesian networks are a natural representation for reasoning about them. In particular, using a BN we can easily calculate:

$$P(e|m) \quad (1)$$

where e is an event (or evidence for an event) and m is the model. However, there is no generally accepted method of classifying an event as anomalous using a BN. Often (e.g., [8,11]), the probability above is tested against a threshold t :

$$P(e|m) < t \rightarrow \text{anomalous} \quad (2)$$

Or, if there is a sequence of events—as there is when trying to detect anomalous *behaviour*—these probabilities may be aggregated over time, as in:

$$\frac{1}{N} \sum_i P(e_i|m) < t \rightarrow \text{anomalous} \quad (3)$$

with N timesteps i . We can choose N either to range over the course of the entire behaviour (i.e., event sequence) or to restrict it to specific time windows.

An alternative approach to identifying anomalies is to check for conflicts within a set of evidence. This is similar to the approach that Das and Schneider [16] take, but within the context of BN inference. Jensen et al. [18] proposed a “conflict measure” to detect possible incoherence in evidence $\mathbf{E} = \{E_1 = e_1, \dots, E_m = e_m\}$:

$$C(\mathbf{E}) = \log \frac{P(E_1 = e_1) \times \dots \times P(E_m = e_m)}{P(\mathbf{E})}$$

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلید کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات