Semi-automated maritime vessel activity detection using hidden Markov models

J du Toit*   JH van Vuuren†

Abstract

Maritime surveillance systems make use of a dearth of sensor data which often include spatio-temporal vessel updates provided by vessels fitted with onboard self-reporting Automatic Identification Systems. These spatio-temporal updates supply low-level information to an operator tasked with observing the surveillance scene and identifying threatening or undesirable behaviour. In this situation, the operator is thus required to interpret the updates by attaching semantic or high-level information to these data.

To this end, automatic activity detection is pursued in this paper as a means to describe vessel motion patterns within the surveillance scene. In particular, the activity of vessels travelling along a well-established route is investigated. Spatial regions of interest are extracted from historical data using a simple spatial clustering technique. The resulting data set is further reduced by removing outliers subject to chosen features before the remaining patterns are clustered. With the assistance of an operator, who may attribute activities to clusters that have some geographical or behavioural meaning, this approach may contribute to a rudimentary understanding of the scene. The motion patterns within these clusters provide the training data for hidden Markov models which are tasked with classifying newly observed motion patterns that engage in the suggested activity. This process of enriching the vessel updates with semantics is expected to lead to more effective decision making on the part of a maritime surveillance operator who may thus direct cognitive resources towards unknown activities.

Key words: Maritime surveillance, motion patterns, activity detection, hidden Markov models, DBSCAN, dynamic time warping, partitioning around medoids.

1   Introduction

The act of surveillance is the systematic observation of regions, entities or objects by visual, aural, electronic or other means [7]. Integrated systems designed for such tasks have been deployed over a broad spectrum of scenarios, typically with the common purpose of providing security. Areas of application include traffic monitoring on motorways, the

---

*Department of Logistics, Stellenbosch University, Private Bag X1, Matieland, 7602, South Africa, email: 13421530@sun.ac.za

†**(Fellow of the Operations Research Society of South Africa)**, Department of Industrial Engineering, Stellenbosch University, Private Bag X1, Matieland, 7602, South Africa, email: vuuren@sun.ac.za

monitoring of public spaces and quality control of industrial processes [6, 8, 14]. As many of these systems mature, their focus invariably expands from observation and data collection to automated decision support and intelligent action[1]. In the video surveillance community, these systems are therefore referred to as second generation systems [20]. This shift to automation assists surveillance operators by alleviating their workload and potentially directing their attention to activities of interest.

These notions are also of import in maritime surveillance, which is driven by the societal demand for improved security. The large volumes of trade passing through ports[2] along with the responsibility of coastal states to protect their natural resources, their citizens, sea-faring vessels and the personal safety of mariners, necessitates such measures. Any maritime entities which endanger these notions are considered *threats* or are considered to be exhibiting threatening behaviour. For example, in a vessel traffic control scenario, the threat of collision is ever present and acts of poaching and pollution are threatening to the environment. The timely identification and effective response to threats is of great importance in mitigating their impact or neutralising them entirely.

In this paper a data-driven approach is taken to identifying the activity of vessels travelling along a well-established route. A similar approach was applied with success in traffic surveillance [13] and its applicability to *Automatic Identification System* (AIS) data is explored in this paper, where the training data are determined using data mining.

2   Related work and discussion

Algorithms capable of identifying activities in a maritime surveillance picture rely in part on spatio-temporal sensory data[3]. The participation in an activity by an entity may be determined from a single datum or from a sub-sequence of its recorded trajectory data (an example of the former is the event that a vessel enters a no-go area, while an example of the latter is the *Williamson turn* manoeuvre which brings a vessel about in the case of a man overboard). However, a significant amount of information is captured in the sequential and temporal nature of the data, requiring time and resource consuming spatio-temporal analyses by sub-components of such a surveillance system. A fundamental problem in time-series analysis and storage is the question of representation. Various approaches towards time series-representation, along with series indexing and similarity, have been researched extensively by the data mining community [4]. Time-series matching has been investigated from both the perspectives of entire sequence matching and sub-sequence matching, using methods such as *dynamic time warping* (DTW) [19] and the *longest common sub-sequence* model [21].

---

[1] *Closed circuit television* camera systems bear testament to this evolution. As technology has advanced and larger quantities of data can be collected, it has become necessary to develop systems capable of automatically detecting events [20].

[2] Approximately 95% of all trade in the Southern African Development Community passes through South African and East African ports [16].

[3] Cooperative reporting mechanisms employed by authorities include the requirement that a vessel docking at a South African port notifies the Maritime Rescue Coordination Centre ninety six hours before arrival at port. The self-reporting *automatic identification system* is another example of a cooperative data source, whereas radars are non-cooperative data sources which collect vessel kinematic data.

The non-trivial task of pattern discovery in temporal data has been approached using distance-based clustering techniques and the statistical model-based technique of *hidden Markov models* (HMMs) [15]. HMMs are also ideally suited as classifiers of temporal sequences where the parameters of an HMM are estimated from training data. This approach has been applied successfully in various contexts where trained HMMs act as representatives of the class on which they were trained [13].

In order for a surveillance system to be able to assist operators in the decision making process of identifying threats at sea, mechanisms are required by which semantics may be added to low-level information associated with vessels, such as raw kinematic data. De Vries *et al.* [2] achieve this by enriching their vessel trajectories with geographical domain knowledge whilst Makris *et al.* [13] identify spatial regions and model the dynamics of objects moving between them using HMMs in a traffic surveillance setting. Although the former method is concerned with identifying high-level behaviours in the maritime domain, many researchers direct their efforts towards the identification of anomalous behaviour in vessel traffic [9, 10, 18]. These approaches are predominantly concerned with instantaneous updates and attempt to build a picture of normality based on historical data. One such approach is the discretization of the surveillance picture into cells on which *kernel density estimation* (KDE) is performed [11]. The data points within those cells are assumed to represent normal vessel behaviour and vessels that deviate significantly from this estimated cell-model are flagged as anomalous. All of these approaches rely on vast quantities of data, and although methods such as KDE are unsupervised, classification techniques such as HMMs additionally require labelled training data. It should also be noted that illicit and threatening behaviour is very seldom observed, thus making their modelling via data-driven approaches more difficult. Anomaly detection thus assumes that all anomalous behaviour is to be considered threatening. As a counterpoint to this approach, rule-based approaches are typically employed to identify specific threats, such as vessels in pursuit of one another or vessels meeting at sea [1]. However, a disadvantage of this approach is that each threatening scenario must then be determined and integrated into the system beforehand.

3 Methodology

A filtering approach is pursued in this paper in combination with clustering in an effort to extract trajectories that are relatively compact in space. Contrary to traffic surveillance applications, vessel trajectories are geographically less constrained. It is expected that trajectories along established routes should be well represented in origin and destination clusters[4]. The method of *density-based spatial clustering of applications with noise*, referred to as DBSCAN, is used as the spatial clustering technique for which it is sufficient to consider simplified incarnations of the trajectories. The poly-line simplification technique of *Douglas-Peucker* (DP) [5] reduces the number of points in a piecewise linear path while retaining its shape. The simplified trajectories are clustered by DBSCAN (using the

---

[4]A decision support system would likely apply viewport clipping or have particular models in operation over certain areas of interest. These viewports would induce origin and destination regions at their boundaries.

same threshold used in DP).

These regions are used to further reduce the data set by discarding trajectories that do not have them in common and trajectories that begin and end in the same region. This reduced data set is expected to contain trajectories that differ markedly from the majority of retained trajectories, and outlier removal with respect to the derived attribute of *sinuosity* is pursued.

Lastly, a final clustering is obtained via the *partitioning around medoids* (PAM) method [3] which utilises *dynamic time warping* (DTW) as a measure of the similarity between vessel trajectories with respect to their positional data. This alignment technique allows contractions or dilations of the temporal axis whilst determining an optimal alignment between two sequences subject to monotonicity and step continuity constraints [3].

These clusters provide the training data for as many HMMs. Sequential data may be represented by an HMM in which it is assumed that the underlying Markov process is hidden, but that the process emits observable quantities. In this setting, the aforementioned features are observable, but the dynamical process leading to their observation is considered unobservable. Suppose the k-th observable feature in a set of M features is denoted by v_k. Then an HMM [17] is specified by a set $S = \{S_1, S_2, \ldots, S_N\}$ of hidden states and an associated $N \times N$ transition probability matrix $A = [a_{ij}]$ that describes the probability of transitioning from state S_i to S_j, *i.e.* $a_{ij} = p(q_{t+1} = S_j | q_t = S_i)$ for all $1 \leq i, j \leq N$, where $a_{ij} \geq 0$ and q_t denotes the state occuring at time t. Furthermore, a probability distribution describing the probability that v_k may be observed when the system is in state S_j is captured in an $N \times M$ *emission* matrix $B = [b_j(k)]$, where $b_j(k) = p(v_k$ at time $t | q_t = S_j)$ for all $1 \leq j \leq N$ and all $1 \leq k \leq M$. Finally, the initial state probability distribution is represented by $\pi = [\pi_i]$, where $\pi_i = p(q_1 = S_i)$ for all $1 \leq i \leq N$.

Once the HMMs have been trained, it is possible to determine whether an observed trajectory was generated by a particular HMM, by determining whether the log-likelihood of the observation, given the HMM model, is greater than a specified threshold.

4   Vessel data

A data set comprising AIS reports was obtained from vessels in the region of Cape Town harbour for use in this paper (see Figure 1(a)). Only the kinematic quantities of the vessels were considered, *i.e.* each vessel's position (reported in geographical coordinates), heading and speed, together with a time-stamp associated with the report. Reports emanating from within a nine-kilometre radius of a chosen reference point at Cape Town harbour over a two-month period were used in the analysis (the reference point is indicated by the filled black circle in Figure 1(a)). Spurious updates were discarded[5] and individual trajectories were partitioned into *stop* and *move* segments. A vessel is deemed to have come to a stop if the reported speed falls below a particular threshold for a sufficient amount of time.

---

[5]Reports have been found to contain reported speeds of hundreds of knots whilst some reported positions would require a vessel to achieve impossible speeds in order to reach that location in the provided time.

(a) AIS data within a $9\,km$ radius of the reference point.

(b) The sample means and sample covariances for the bivariate positional data.
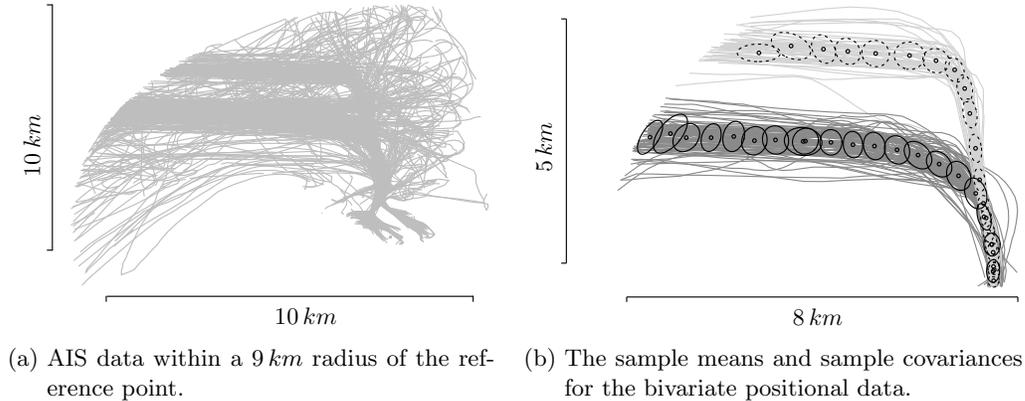
Figure 1: The data used for training and testing.

A clustering of vessel data into two clusters is shown in Figure 1(b) in which the the upper cluster is composed of vessel trajectories resulting from vessels departing from Cape Town harbour whilst the lower cluster comprises trajectories entering the scene and travelling to the harbour. The number of vessel trajectories utilised in this paper are presented in Table 1.

	Total trajectories	Training set	Validation set	Test set
Entry cluster	97	67	10	20
Exit cluster	53	30	8	15
Unclustered	280			

Table 1: The number of vessel trajectories used in this paper.

5   HMM Structure and Training

A *left-to-right* HMM structure was utilised in capturing the sequential nature of the vessel reports. These reports are described by a state-dependent distribution $p(X_t|S_t)$, where $X_t = (x_t, y_t, u_t, v_t)$ is an observation vector comprising positional (x_t, y_t) and velocity (u_t, v_t) data at time t. This distribution was modelled as a multivariate Gaussian distribution, in keeping with the approach of [13]. Furthermore, the *Bayesian information criterion*[6] was used to inform the choice, in favour of mixture distributions.

The initial transition probability matrix A was specified as an upper triangular matrix ($a_{ij} = 0$ for all $i > j$) so as to ensure that the left-to-right structure is maintained during parameter estimation and to lessen the number of free parameters. Similarly, for each HMM, the initial state distribution was taken as $\pi_0 = (1, 0, \ldots, 0)$, thus ensuring that each trajectory begins in the first state. The number of states were arbitrarily chosen in this instance, but they may also have been determined via model selection methods

---

[6]The *Bayesian information criterion* is a model selection criterion which is expressed as $BIC = -2\log L + p\log T$, where L is the log-likelihood of the fitted model, p is the number of parameters of the model and T is the number of observations [12]. This measure typically favours models with fewer parameters.

(such as the aforementioned BIC). An HMM was fitted to the training data using the *Baum-Welch* method[7], the convergence of which is assisted by the initialisation of the parameters of the state-dependent distributions. The medoids determined by PAM in the data mining step were linearly interpolated at intervals that produce the desired number of states and all points within a radius of this interval length were used to compute the sample means and covariances for each state-dependent distribution (the bivariate case is illustrated in Figure 1(b)).

6 Activity Classification

The estimated HMMs were considered to represent an activity in the scene and they may be labelled as such. For instance, a great deal of vessel traffic is observed travelling to and from Cape Town harbour that may successfully be classified by an HMM corresponding to the route that they have taken. Thresholds for class membership were determined from a validation set $\mathbb{V}$ (a portion of the test set was withheld during training) by selecting the minimum log likelihood value attained by members of each class with respect to its corresponding HMM. A vessel trajectory $X = X_0, \ldots, X_n$ was deemed to be a member of a class $\mathcal{C}_i$ if the probability of the sequence of observations, given the corresponding HMM λ_i, is less than the selected threshold for the i-th class. That is, $\log p(X|\lambda_i) > \min_{Y \in \mathbb{V}} \log(p(Y|\lambda_i))$. Trajectories that do not feature a displacement of more than five kilometres were not considered for classification.

7 Results

Two HMMs were estimated from the clustered data, corresponding to entry and exit trajectories. The test sets were classified once the thresholds had been chosen for each HMM. The resulting number of *true positive* (TP), *false positive* (FP), *true negative* (TN) and *false negative* (FN) classifications for the test sets are shown in Table 2. The results of calculating the membership of the trajectories in the unclustered data set are also reported in this table.

	Test sets				Unclustered set	
	TP	TN	FP	FN	TN	FP
Entry HMM	22	15	0	0	241	39
Exit HMM	13	22	0	2	278	2

Table 2: The classification results of the entry and exit test sets for each HMM, as well as classification results for the unclustered set.

Closer inspection of the data revealed that vessels approaching the harbour often reduce speed or perform loop manoeuvres in the entry channel. The latter trajectories were removed from the training data set in the data mining phase as they are considered to be outliers with regard to the sinuosity measure, whilst the former are regarded to have

---

[7]The Baum-Welch method is a special case of the *expectation-maximisation* algorithm and is not guaranteed to find a global optimum [12].
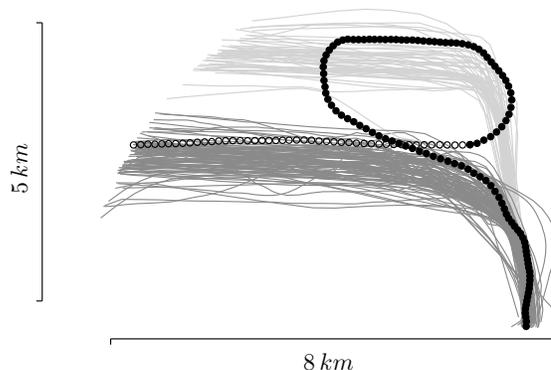
Figure 2: The filled circles indicate the updates for which the entry HMM successfully classified the trajectory.

come to a stop and the trajectories are therefore subdivided into two segments. The data mining process assigned these trajectories to the unclustered set, but the entry HMM still successfully classified them.

Lastly, the responsiveness to a change in behaviour whilst a vessel is underway was investigated. The vessel in Figure 2 changes lanes by travelling along the entry route and then switching to the exit route. The updates along its trajectory for which it was classified by the entry HMM are indicated by open circles. The updates for which the trajectory is no longer described by either of the HMMs, are indicated by solid circles.

8　Conclusion

Selecting position and velocity as the features for training HMMs results in a satisfactory classification of trajectories on the limited AIS test data. However, the use of a single classifier to describe vessels travelling towards the port is insufficient as there are vessels that slow down on their approach and vessels that do not.

The performance of the classification models was found to rely heavily on the training data extracted via the filtering and clustering method. The use of the filters and features of sinuosity were demonstrated to be useful in extracting spatially compact trajectories and the resulting data were shown to lend themselves to HMM classification and training.

References

[1]　BRAX C & NIKLASSON L, 2009, *Enhanced situational awareness in the maritime domain: An agent-based approach for situation management*, Proceedings of the SPIE: Intelligent Sensing, Situation Management, Impact Assessment, and Cyber-Sensing, Orlando (FL), pp. 3–13.

[2]　DE VRIES G, VAN HAGE W & VAN SOMEREN M, 2010, *Comparing vessel trajectories using geographical domain knowledge and alignments*, Proceedings of the 10[th] IEEE International Conference on Data Mining Workshops, Sydney, pp. 209–216.

[3] DU TOIT J & VAN VUUREN JH, 2012, *Towards coastal threat evaluation decision support*, Proceedings of the 41st Annual Conference of the Operations Research Society of South Africa, pp. 31–39.

[4] FU T, 2011, *A review on time series data mining*, Engineering Applications of Artificial Intelligence, **24(1)**, pp. 164–181.

[5] HECKBERT PS & GARLAND M, 1997, *Survey of polygonal surface simplification algorithms*, (Unpublished) Technical Report, Carnegie-Mellon University, Pittsburgh (PA).

[6] HU W, TAN T, WANG L & MAYBANK S, 2004, *A survey on visual surveillance of object motion and behaviors*, IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews, **34(3)**, pp. 334–352.

[7] INCE AN, TOPUZ E & PANAYIRCI E, 2000, *Principles of integrated maritime surveillance systems*, Kluwer Academic, Norwell (MA).

[8] KUMAR A, 2008, *Computer-vision-based fabric defect detection: A survey*, IEEE Transactions on Industrial Electronics, **55(1)**, pp. 348–363.

[9] LANE R, NEVELL D, HAYWARD S & BEANEY T, 2010, *Maritime anomaly detection and threat assessment*, Proceedings of the 13th International Conference on Information Fusion (FUSION'10), pp. 1–8.

[10] LAXHAMMAR R, 2008, *Anomaly detection for sea surveillance*, Proceedings of the 11th International Conference on Information Fusion (FUSION'08), pp. 55–62.

[11] LAXHAMMAR R, FALKMAN G & SVIESTINS E, 2009, *Anomaly detection in sea traffic—A comparison of the Gaussian mixture model and the kernel density estimator*, Proceedings of the 12th International Conference on Information Fusion (FUSION'09), pp. 756–763.

[12] MACDONALD IL & ZUCCHINI W, 2009, *Hidden Markov and other models for discrete valued time series*, CRC Press, Boca Raton (FL).

[13] MAKRIS D & ELLIS T, 2005, *Learning semantic scene models from observing activity in visual surveillance*, IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics, **35(3)**, pp. 397–408.

[14] MCCAHILL M & NORRIS C, 2002, *CCTV in Britain*, (Unpublished) Technical Report, Centre for Criminology and Criminal Justice, University of Hull, Hull.

[15] OATES T, FIROIU L & COHEN P, 1999, *Clustering time series with hidden Markov models and dynamic time warping*, Proceedings of the IJCAI-99 Workshop on Neural, Symbolic and Reinforcement Learning Methods for Sequence Learning, pp. 17–21.

[16] *Ports and ships*, [Online], [Cited June 10[th], 2013], Available from `http://ports.co.za/index.php`.

[17] RABINER LR, 1989, *A tutorial on hidden Markov models and selected applications in speech recognition*, Proceedings of the IEEE, **77(2)**, pp. 257–286.

[18] RISTIC B, LA SCALA B, MORELANDE M & GORDON N, 2008, *Statistical analysis of motion patterns in AIS data: Anomaly detection and motion prediction*, Proceedings of the 11th International Conference on Information Fusion (FUSION'08), pp. 40–46.

[19] SALVADOR S & CHAN P, 2007, *Toward accurate dynamic time warping in linear time and space*, Intelligent Data Analysis, **11(5)**, pp. 561–580.

[20] VALERA M & VELASTIN S, 2005, *Intelligent distributed surveillance systems: A review*, Proceedings of the 2[th] IEEE Conference on Vision, Image and Signal Processing, pp. 192–204.

[21] VLACHOS M, KOLLIOS G & GUNOPULOS D, 2002, *Discovering similar multidimensional trajectories*, Proceedings of the 18th International Conference on Data Engineering, San Jose (CA), pp. 673–684.