Modelling massive AIS streams with quad trees and Gaussian Mixtures

Anita Graser
AIT Austrian Institute of
Technology
Giefinggasse 2
1210 Vienna, Austria
anita.graser@ait.ac.at

Peter Widhalm
AIT Austrian Institute of
Technology
Giefinggasse 2
1210 Vienna, Austria
peter.widhalm@ait.ac.at

Abstract

Pressing issues related to the movement of people and goods can be tackled today thanks to improvements in tracking and communications technology that have made it possible to collect movement data on a big scale. Maritime data from the Automatic Identification System (AIS) is one of the fast growing sources of movement data. Existing approaches for AIS data analysis suffer from scalability issues. Therefore, scalable distributed modelling and analysis approaches are needed. This paper presents a novel scalable movement data model that takes advantage of an adaptive grid based on quad trees. Our data model supports anomaly detection in massive movement data streams by combining advantages of both grid and vector-based approaches. We demonstrate the applicability of this approach for anomaly detection in AIS datasets comprising 560 million location records.

Keywords: AIS, movement data, trajectories, data models, maritime, quad tree

1 Introduction

Movement of people and goods is related to many of the most pressing issues we are facing today. Emissions from the transport sector contribute significantly to climate change and casualties from transport accidents are rising quickly, particularly in developing countries with strong population growth and high urbanization rates. In the context of maritime movement data, consolidation of control centres on the one hand and increasing vessel traffic on the other hand pose challenges to the environment as well as traffic safety. One of the fastest growing data sources is the Automatic Identification System (AIS) which is mandatory for ships above a certain size (SOLAS, 2002; Commission of the European Communities, 2008). AIS is a cooperative tracking system, meaning that vessels actively broadcast their location, status, vessel, and voyage information. The reporting interval between location updates depends on the vessel speed and ranges from 2 seconds (at high speeds while changing course) to 3 minutes (while anchored). Live AIS data is used, for example, to monitor vessel movement in vessel traffic services (VTS) operated by port authorities. Historical AIS data is used for planning purposes, for example, to evaluate traffic separation schemes (TSS) which define shipping routes. Historically, monitoring systems used by coastal authorities are similar to systems used on board of individual vessels. These systems were never designed to deal with larger amounts of data. The goal of novel AIS monitoring systems is to help decision makers to deal with the data load by reducing clutter and guiding operator attention. New approaches aim to summarize "normal traffic state" and its variations in order to detect unusual events. To goal of automatic anomaly detection is to enable effective maritime situational awareness in order to support the operator of maritime surveillance systems in detecting suspicious situations or threats and in taking appropriate action (Sidibé & Shu, 2017). To this end, different aggregation methods and models for massive movement data are being developed

Common data models for anomaly detection (in movement data in general and AIS data in particular) use either a grid-based (Laxhammar et al., 2009) or a vector-based (Vespe et al., 2012; Palotta et al., 2013) approach. Many vector-based data models have been developed to avoid the challenges of selecting an optimal cell size for grid-based approaches as well as the computational burden resulting from increasing grid resolution (Palotta et al., 2013). Regardless of these efforts though, current approaches suffer from scalability issues (Sidibé & Shu, 2017). The volume of AIS data (historic AIS data will quickly grow into TB ranges) requires distributed data mining algorithms (Dobrkovic et al., 2016b). In the context of massive movement data, grids provide the clear advantage that they already partition data into cells that can be distributed over multiple computing nodes in a cluster.

With this massive movement data scalability issue in mind, we present a novel movement data model combining the advantages of grid and vector approaches. The challenge of optimal grid cell size selection versus computational burden is dealt with by generating an adaptive grid using a quad tree approach. We demonstrate the usefulness of this approach for the anomaly detection in AIS data use case. Anomalies are detected on an individual AIS record level. This paper is organized as follows: Section 2 presents related work on quad tree data models for movement data. Section 3 introduces our methodology and presents a real life example. Finally, Section 4 provides an outlook on future work.

2 Quad trees for movement data

One of the main challenges of modelling movement data in a grid is to pick an optimal cell size which provides a good balance between necessary details on the one hand and computational burden of increasing resolution on the other

(Palotta et al., 2013). In the context of modelling AIS data, for example, Osekowska et al. (2014) find that optimal grid cell size differs between regions: within harbour areas, they identify an optimal cell size of 60 to 200 meters, while on the open sea optimal cell size ranges from 300 to 1,000 meters.

Quad trees provide a way to divide space into a more flexible grid that combines smaller and bigger cells as needed. Quad trees are commonly used to index data in two-dimensional space, and can thus be used to index two-dimensional spatial data. In the context of movement data analysis, Wang et al. (2008) suggest a quad tree combined with a temporal index for trajectories. More recently, Xie et al. (2016) use an oct-tree, a three-dimensional version of a quad-tree, to store trajectory locations in their big data structure. Besides indexing, Zhao (2015) uses a quad tree to partition urban space based on taxi trajectories, and Ho & Ruan (2011) use a quad tree to discover interesting locations for their spatial data privacy solution.

Papers dealing with AIS analysis and quad trees include Dobrkovic et al. (2016a) who use a quad tree to identify areas of high AIS data density which they feed into a genetic algorithm that identifies typical waypoints. On a different note, Xu et al. (2016) use a quad tree to model AIS signal reception rate with different antenna. The work most closely related to our approach is by Woxberg & Grahn (2015) who propose dynamic potential field grids using quad tree structures to detect anomalous vessel positions. In contrast, our approach goes beyond just modelling vessel locations on an individual record level. Besides location, our data model also deals with additional ordinal and categorical data. For our AIS anomaly detection use case, this additional information includes vessel speed, heading, type, and status.

The majority of proposed approaches for modelling movement data from AIS use either Kernel Density Estimation (KDE) or Gaussian Mixture Models (GMM) (Tu et al., 2017; Sidibé & Shu, 2017). Most proposed clustering approaches are based on different versions of the DBSCAN algorithm. However, DBSCAN's complexity is $O(n^2)$ or $O(n \log n)$ if spatial indexes were used which hasn't been the case in recent work (Sidibé & Shu, 2017). To the best of our knowledge, there currently exists no work describing a scalable movement data model that supports anomaly detection. To this end, our approach combines grid-based and vector-based approaches: quad trees for spatial data partitioning, as well as multiple prototypes per grid cell with GMM to describe local movement characteristics.

3   Methodology

AIS records contain information about the current motion state of a vessel, in particular: position, speed over ground, heading, course over ground, and rate of turn. We model the multivariate distribution of the motion state vector as a 6-dimensional Gaussian Mixture distribution. The model is used detect unusual events by comparing newly arriving AIS records to the model. As for model fitting, instead of using standard Expectation Maximization (EM) to fit the Gaussian Mixture Model (GMM), we employ an incremental approximate algorithm similar to Vector Quantization (Gray, 1984) and Leader-Follower Clustering (Duda et al, 2001). In contrast to EM which works on the entire data set to perform model fitting, our algorithm allows processing large historical data sets sequentially since Vector Quantization allows iterative updates of mixture components ("prototype vectors" in Vector Quantization) and the Leader-Follower approach supports online clustering without a predefined number of clusters and better run time performance than DBSCAN. Our algorithm thus can continuously update the model using potentially endless streams of new input data.

The mixture components are organized in a dynamically growing quad tree which partitions the modelled area into a hierarchical spatial grid as shown in the example in Figure 1. Grid cells and the corresponding tree branches are added as needed to represent the data. Each grid cell is represented by a variable number of Gaussian components, depending on the number of observations in that cell, their distribution and a predefined maximum number of components in each cell.

The tree structure provides the following advantages:
- When a new AIS record is added, the relevant components can be efficiently accessed to update the model
- Model components in different branches of the tree can be updated independently, allowing easy parallelization of the algorithm, and
- Mixture components can be efficiently aggregated at different levels up the tree hierarchy for representations at different levels of detail (e.g. zoom-levels).

Figure 1: Grid of cells with at least 1,000 records in the observation period (8 weeks in 2015 and 2016).

The complete grid underlying the example shown in Figure 1 (which is filtered to only show grid cells with at least 1,000 records) consists of 369,153 cells on zoom levels 0 to 13, where zoom level 0 is a square of extent -180/-180/180/180 degrees and each subsequent level splits each cell into four. A complete grid at zoom level 13 would therefore consist of $4^{13} = 67,108,864$ cells. This illustrates the memory saving potential of our approach.

3.1   Algorithm description

The complete algorithm for a quad tree with GMM of depth L can be summarized as follows:
- Start with an empty data structure representing the initially empty quad tree
- For each new AIS record r, compute the spatially matching node at each level of the quad tree from its root to the node at level L

- Create all missing nodes and branches in the tree structure
- Update the Gaussian mixture distribution of the matching node N_L^* at level L:
  - Let C be the set of Gaussian components of N_L^*
  - Find the most similar prototype vector, i.e. the mean vector of a Gaussian component

$$\mu^* = \arg \min_{c \epsilon C} \|r - \mu_c\| \qquad (1)$$

  - If $\|r - \mu^*\| < d_{max}$:
    - Update mean vector μ^*, covariances S^*, and number of samples from which the Gaussian component was estimated
  - Else:
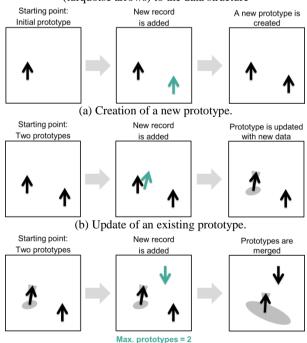   - Add a new Gaussian component with $\mu_{new} = r$ and $S_{new} = 0$
    - If $|C| > n_{max}$: find and merge most similar prototype vectors

$$(\mu_x, \mu_y) = \arg \min_{x \epsilon C, y \epsilon C} \|\mu_x - \mu_y\| \qquad (2)$$

In the algorithm above, n_{max} is the maximum number of components in each cell and d_{max} is a threshold for the distance between the new record and the most similar prototype vector.

Figure 2 illustrates the three different concepts for creating (a), updating (b), and merging (c) prototype vectors.

Figure 2: Illustration of the process of adding records (turquoise arrows) to the data structure

(a) Creation of a new prototype.

(b) Update of an existing prototype.

(c) Merging of existing prototypes due to necessity of creating a new prototype) that would exceed maximum allowed number of prototypes.
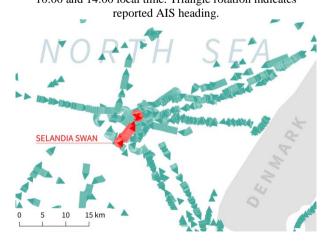
3.2 Anomaly detection

After an initial calibration phase, the model is ready to detect unusual events by comparing newly arriving AIS records to the model. One downside of the grid based organization of the model are possible discontinuities in the modelled distribution along the borders between neighboring cells. Therefore, AIS records are not only compared to the Gaussian components of the matching cell, but also to Gaussian components in neighboring cells. If no spatially matching cell exists, the record is considered unusual as there is no past observation the new record could be compared to. Otherwise, we select the Gaussian component c^* with minimal Mahalanobis distance d_M to the AIS record x

$$c^* = \arg \min_c d_M(x, c) = \sqrt{(x - \mu_c)^T S_c^{-1}(x - \mu_c)} \qquad (3)$$

and perform a chi-squared test to assess the goodness of fit of the new AIS record to the modelled distribution. Records with bad model fit are also reported as unusual events. Examining the deviation between record and model along each dimension of the motion state vector allows to identify reasons for the unusual events, such as travelling with too high speed or unusual course.

Figure 3 provides an example of anomalies that were detected in the context of an actual marine accident, a fall over board from SELANDIA SWAN on 23 July 2015 (DMAIB, 2016). The 3rd officer fell overboard between 10:05 and 10:10 local time. Within 15 minutes the crew initiated man overboard procedures. Several ships participated in the search. Our approach correctly detects numerous anomalous records of SELANDIA SWAN (shown in red) as well as other ships (shown in turquoise). The main reasons why these records haven been flagged as anomalous are heading (41%), slow speed (25%), and position (14%) due to the unusual patterns observed while the vessels were searching for the missing officer.

Figure 3: Detected anomalous records during the marine incident of SELANDIA SWAN in 2015. Anomalous records of Selandia Swan (red) and other ships (turquoise) between 10:00 and 14:00 local time. Triangle rotation indicates reported AIS heading.

4   Conclusion

We demonstrated a novel modeling approach combining the advantages of grid and vector based data models for anomaly detection in the context of massive movement data. The challenge of optimal cell size selection versus computational burden is dealt with by generating an adaptive grid using a quad tree approach. Anomalies are detected on an individual record level when a movement does not conform to the modelled distribution.

We demonstrated the applicability of this approach for detecting anomalies in AIS data with an example of an actual incident. However, to evaluate the performance and practical usefulness of this approach in real-life settings it will be necessary to involve human domain experts to collect feedback about the actual significance of the detected anomalies.

Future work will expand this approach to detect unusual traffic flows, when each individual movement on its own is normal, but overall, for example, there is much more movement in an area than usual.

Acknowledgements

References

Commission of the European Communities. DocumentCOM(2008) 310 final - 2005/0239 COD, Brussels11-06-2008. [Online] Available from: http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=COM:2008:0310:FIN [Accessed 20th December 2017].

Danish Maritime Accident Investigation Board (DMAIB) (2016) Marine accident report about fall over board from SELANDIA SWAN on 23 July 2015. [Online] Available from: http://www.dmaib.com/News/Sider/Marineaccidentreportaboutfalloverboardon23July2015.aspx [Accessed 20th December 2017].

Dobrkovic, A., Iacob, M. E., & Van Hillegersberg, J. (2016a). Maritime Pattern Extraction from AIS Data Using a Genetic Algorithm. In Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference on (pp. 642-651). IEEE.

Dobrkovic, A., Iacob, M. E., van Hillegersberg, J., Mes, M. R., & Glandrup, M. (2016b). Towards an approach for long term AIS-based prediction of vessel arrival times. In Logistics and Supply Chain Innovation (pp. 281-294). Springer International Publishing.

Duda, R. O., Hart, P. E., and Stork, D. G. (2001), Pattern Classification. Wiley-Interscience, p. 654.

Gray, R. (1984). "Vector quantization," in IEEE ASSP Magazine, vol. 1, no. 2, pp. 4-29. doi: 10.1109/MASSP.1984.1162229

Ho, S. S., & Ruan, S. (2011). Differential privacy for location pattern mining. In Proceedings of the 4th ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS (pp. 17-24). ACM.

Laxhammar, R., Falkman, G., & Sviestins, E. (2009). Anomaly detection in sea traffic-a comparison of the gaussian mixture model and the kernel density estimator. In Information Fusion, 2009. FUSION'09. 12th International Conference on (pp. 756-763). IEEE.

Osekowska, E., Johnson, H., & Carlsson, B. (2014). Grid size optimization for potential field based maritime anomaly detection. Transportation Research Procedia, 3, 720-729.

Pallotta, G., Vespe, M., & Bryan, K. (2013). Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction. Entropy, 15(6), 2218-2245.

Sidibé, A., & Shu, G. (2017). Study of Automatic Anomalous Behaviour Detection Techniques for Maritime Vessels. The Journal of Navigation, 70(4), 847-858.

Safety of Life at Sea (SOLAS). (2002). convention Chapter V.Regulation 19. [Online] Available from: http://solasv.mcga.gov.uk/regulations/regulation19.htm. [Accessed 20th December 2017].

Tu, E., Zhang, G., Rachmawati, L., Rajabally, E., and Huang, G. B. (2017). Exploiting AIS Data for Intelligent Maritime Navigation: A Comprehensive Survey From Data to Methodology. IEEE Transactions on Intelligent Transportation Systems, PP, no. 99, pp. 1-24.

Vespe, M., Visentini, I., Bryan, K., & Braca, P. (2012). Unsupervised learning of maritime traffic patterns for anomaly detection. NATO Undersea Research Centre, La Spezia, Italy.

Wang, L., Zheng, Y., Xie, X., & Ma, W. Y. (2008). A flexible spatio-temporal indexing scheme for large-scale GPS track retrieval. In Mobile Data Management, 2008. MDM'08. 9th International Conference on (pp. 1-8). IEEE.

Woxberg, L., & Grahn, S. (2015). Maritime anomaly detection with dynamic potential field grids. Master thesis. Blekinge Institute of Technology, Faculty of Computing. Sweden. [Online] Available from: http://www.diva-portal.org/smash/record.jsf?pid=diva2%3A839528&dswid=-8993 [Accessed 20th December 2017].

Xie, X., Mei, B., Chen, J., Du, X., & Jensen, C. S. (2016). Elite: an elastic infrastructure for big spatiotemporal trajectories. The VLDB Journal, 25(4), 473-493.

5

Xu, T., Hu, Q., Xiang, Z., Yang, C., & Wang, D. (2016). The Comparison Study on AIS Signal Reception Rate with Directional Antenna and Omni Antenna. TransNav: International Journal on Marine Navigation and Safety of Sea Transportation, 10.

Zhao, K. (2015). Understanding urban human mobility for network applications. PhD thesis. Department of Computer Sciences, University of Helsinki, Finland.