Challenges in Vessel Behavior and Anomaly Detection: From Classical Machine Learning to Deep Learning*

Lucas May Petry[1], Amilcar Soares[2], Vania Bogorny[1], Bruno Brandoli[2], and Stan Matwin[2]

[1] Universidade Federal de Santa Catarina (PPGCC/UFSC), Florianópolis, Brazil
lucas.petry@posgrad.ufsc.br
[2] Institute for Big Data Analytics, Dalhousie University, Halifax, Canada

Abstract. The global expansion of maritime activities and the development of the Automatic Identification System (AIS) have driven the advances in maritime monitoring systems in the last decade. Given the enormous volume of vessel data continuously being generated, real-time analysis of vessel behaviors is only possible because of decision support systems provided with event and anomaly detection methods. However, current works on vessel event detection are ad-hoc methods able to handle only a single or a few predefined types of vessel behavior. Most of the existing approaches do not learn from the data and require the definition of queries and rules for describing each behavior. In this paper, we discuss challenges and opportunities in classical machine learning and deep learning for vessel event and anomaly detection.

Keywords: Automatic identification system · Behavior detection · Anomaly detection · Spatiotemporal data mining.

NOTICE: This is an accepted article to be published by Springer in the proceedings of the 33rd Canadian Conference on Artificial Intelligence.

1 Introduction

The worldwide growth of maritime traffic and the development of the Automatic Identification System (AIS) has led to advances in monitoring systems for

* This study was financed in part by the Brazilian agencies CNPq and CAPES - Project Big Data Analytics: Lançando Luz do Genes ao Cosmos (CAPES/PRINT process number 88887.310782/2018-00), Fundação de Amparo a Pesquisa e Inovação do Estado de Santa Catarina - Project Match (co-financing of H2020 Projects - Grant 2018TR 1266), and the European Union's Horizon 2020 research and innovation programme under Grant Agreement 777695 (MASTER). The authors also acknowledge the support of the Natural Sciences and Engineering Research Council of Canada and of Global Affairs Canada - ELAP for this research.

preventing vessel accidents and detecting illegal activities. In addition, the integration of vessel traffic data with environmental and climatological data allows more complex analyses and a better understanding of the cause and effect of maritime events [18]. While preventing vessel accidents means saving money for shipping companies, from the environmental point of view, it also protects the marine fauna and flora from irreversible damage [2]. However, real-time monitoring and analysis of vessel traffic can be overwhelming to maritime agents due to the high volume of data continuously generated. Therefore, decision support systems are fundamental to enabling efficient and effective maritime control.

The detection of events from AIS data has been the subject of study of several works in the literature [16,18,9]. In particular, some approaches have been proposed for detecting events such as changes in the speed [18,19] or in the course of vessels [19], proximity of vessels to other vessels [9], illegal fishing [16] or possibly hazardous activity [18], among others. However, to the best of our knowledge, most current works are ad-hoc approaches that do not learn from the data and are limited to detecting a restricted set of predefined events. Such methods are not able to detect unforeseen events and also require the assistance of domain specialists for defining rules that characterize each event.

Another aspect often ignored by previous research is the integration of data from different sources for analyzing vessel behavior. Even though this can be advantageous to maritime systems, only a handful of works have addressed it [18]. For example, detecting small vessels heading towards ice-infested waters and that are not equipped for handling this situation allows the decision-maker to warn the captain in advance. Such strategy avoids the deployment of a search and rescue mission, which might represent a high cost (e.g., lives, resources) to maritime authorities.

In this work, we present research gaps and challenges in machine learning for detecting different types of vessel behavior, considering several constraints imposed by real-time data streams and the maritime monitoring domain. We highlight the potential of exploiting machine learning techniques for maritime monitoring, as it has been shown to be fundamental for enabling cognitive smart cities, for instance, which is a scenario similar to ours with heterogeneous sensor data that requires real-time decision making systems [13].

2 Research Challenges and Opportunities

In this section, we explore the challenges and research gaps that are currently faced by vessel behavior monitoring techniques. We address three main tasks related to vessel behavior detection: (i) the actual detection of different vessel behaviors or behavior changes; (ii) identifying and relating recurrent behaviors; (iii) providing the user with means for interpreting and analyzing the detected vessel behaviors. In addition, we discuss about two main issues associated to the data that can be challenging for machine learning methods: (i) the data has limited or no class labels about vessel behaviors; and (ii) often, there is no knowledge whatsoever of what are the behaviors or labels present in the data.

2.1 Behavior Detection: Supervised vs. Unsupervised

Most of the existing works for behavior detection in vessel traffic monitoring do not learn from the data. Instead, algorithms are proposed to detect a restricted set of vessel activities via rules and thresholds defined beforehand. Defining such rules requires the knowledge from domain specialists, which may introduce considerable human bias in the analysis. Although the sole use of supervised learning techniques may seem appealing, they can also limit the detection of events to a predefined set, i.e., the classification labels. With that in mind, we believe that research on unsupervised learning techniques is promising for taking vessel event detection research to the next level.

A closely-related research topic that has already been explored in the literature is the detection of concept drift or change in time series data [1,17]. In streaming data, concept drift is commonly referred to as the detection of significant changes in the data distribution [8]. Several works for change detection in data streams were designed to work along with a supervised learning model, detecting drifts based on the error rate of the learner [1].

Unsupervised approaches have also been proposed for detecting concept drift [1,8,17]. However, they also exhibit some drawbacks, such as being limited to univariate data [1]; simply detecting changes based on individual feature correlation [8] instead of analyzing dependency relationships between features and how they characterize different behaviors; and even lacking direct interpretability of the detected changes [17]. In spite of their limitations, these methods can provide a solid starting point for future works on unsupervised behavior change detection and posterior event notification.

2.2 Identifying Recurrent Behavior Patterns

Detecting multiple instances of the same behavior is an essential factor in maritime monitoring. Besides avoiding multiple analyses of the same behavior by agents, it allows agents to have a higher-level picture of behavior patterns of a single vessel or even of a group of vessels in a region. In previous works for vessel event detection, identifying multiple occurrences of the same behavior was a trivial task, since the detection process consisted mostly of a query [18] or an algorithm for detecting a single behavior [9]. On the other hand, in the unsupervised setting, a pattern or behavior needs to be characterized in a way that multiple occurrences of the same pattern can be identified. Although concept drift or behavior change detection methods can indicate boundaries for different types of behavior in streaming data, they do not provide a way for detecting recurrent behavior.

To the best of our knowledge, the only unsupervised learning method proposed for detecting different and recurrent behaviors in time series data is the Toeplitz Inverse Covariance-based Clustering (TICC), introduced by Hallac et al. [6]. TICC segments multivariate sensor data into sequences of states or clusters (i.e., behavior patterns), representing each behavior as a Markov Random Field (MRF). However, the number of clusters is fixed and must be defined by the user, meaning that the number of different behaviors present in the data should be known a priori. Additionally, TICC is not directly suitable for streaming data, as it assumes that all data is available at the same time, and it requires a few iterations of the algorithm for convergence. Although TICC has some drawbacks, we believe that future research could take advantage of the method for recurrent behavior detection as, for instance, use MRFs to represent and identify the same behavior in different trajectory segments.

2.3 Towards Interpretable Behavior Patterns

Doshi & Kim [4] define interpretability in the context of machine learning as "the ability to explain or to present in understandable terms to a human." In order to enable maritime agents to make data-driven decisions about suspicious or dangerous vessel activity, discovered behavior patterns must be interpretable, i.e., understandable to the agents. Moreover, interpretability may assist in the detection of recurrent behaviors if they can be explicitly characterized based on feature observations.

Interpretability is, perhaps, the main advantage of a few existing works, since rules explicitly define events related to vessel behavior. In contrast to these approaches, detecting interpretable patterns can be challenging for unsupervised machine learning algorithms. For vessel behavior detection, we conjecture that the interpretability of the model may be negligible if guarantees are given about the characteristics of the detected behaviors. For instance, if a behavior can be represented as an MRF, and it is always defined by the same variable dependency graph, the user might not be interested in the details of the underlying model as long as this structure is guaranteed for all future occurrences of the same behavior. Thus, detection methods that are based on abstractions of the real observed features (e.g., [17,12]) can be a feasible option. Afterwards, other techniques could be exploited for correlating and providing interpretability of different behaviors from the real observed variables.

2.4 Diving into Deep Learning

Deep learning models have often been set aside in favor of linear models, because of the claimed lack of interpretability that they have [10].We believe that for a similar reason only very few works have addressed anomaly detection in the maritime domain with deep learning [14,20]. To the best of our knowledge, no work has addressed the detection of specific vessel behaviors with deep learning techniques. More recently, however, a few works have been proposed to assist in the visual interpretation of these models (e.g. [21]), while others have even questioned previous claims over the interpretability of deep models [10].

Computer vision has experienced the most advances with deep learning, and Convolutional Neural Networks (CNNs) are nowadays widely used for image classification [7]. An intuitive approach is the use of satellite imagery data with CNNs for detecting ship anomalies. However, it may be difficult or even impossible to detect certain vessel behaviors from satellite imagery, which limits the use of modern deep learning methods exclusively based on image data [15]. Also, obtaining satellite images is generally more expensive in comparison with AIS data. On the other hand, CNNs have already been used for trajectory classification [3] and prediction [11], based on trajectory features that can be extracted from AIS data. Therefore, CNNs with visual techniques for interpretability [21] could be further exploited for behavior and anomaly detection.

2.5 Big, Yet Limited Data

The growth of maritime activity led to advances in AIS technology, and such developments resulted in large volumes of data being generated every day. Although a large amount of AIS data is available, it lacks labels. Labels are a valuable piece of information for researchers and machine learning algorithms. However, labeling data is difficult and expensive since it requires knowledge from domain experts.

The lack of labels has guided the focus of current research either to detecting a single type of vessel behavior in an unsupervised manner [9], or to proposing the detection of different events via predefined behavior rules, materialized in the form of data queries [16,18]. Future works could take advantage of the knowledge described in previous approaches for labeling data, to provide input-output examples to machine learning algorithms, as well as ground truth for evaluating novel approaches. On the other hand, research could concentrate on synthesizing new behavior data with Generative Adversarial Networks [5], for example, for enhancing the performance of other supervised models.

3 Summary and Final Remarks

Even though maritime monitoring has experienced significant progress in the last decade, most of existing works do not take full advantage of machine learning techniques for vessel behavior detection. In fact, existing research has focused on proposing queries, rules, and ad-hoc algorithms for detecting specific types of behavior. We argue that this methodology inhibits further advances to more general behavior detection approaches, constraining monitoring systems to function only under frequent human supervision. In this paper, we presented research gaps in the field, indicating points of improvement and opportunities for future works. We hope to instigate the development of new algorithms, methods, and tools for ship behavior monitoring, as many aspects of it still remain unaddressed.

References

1. Bifet, A., Gavalda, R.: Learning from time-changing data with adaptive windowing. In: Proc. of the 2007 SIAM int. conf. on data mining. pp. 443–448. SIAM (2007)
2. Claramunt, C., Ray, C., Salmon, L., Camossi, E., Hadzagic, M., Jousselme, A., Andrienko, G., Andrienko, N., Theodoridis, Y., Vouros, G.: Maritime data integration and analysis: recent progress and research challenges. Adv. in Database Technology - EDBT **2017**, 192–197 (2017)

3. Dabiri, S., Heaslip, K.: Inferring transportation modes from gps trajectories using a convolutional neural network. Transportation research part C: emerging technologies **86**, 360–371 (2018)
4. Doshi-Velez, F., Kim, B.: Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608 (2017)
5. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Adv. in neural inf. processing systems. pp. 2672–2680 (2014)
6. Hallac, D., Vare, S., Boyd, S., Leskovec, J.: Toeplitz inverse covariance-based clustering of multivariate time series data. In: Proc. of the 23rd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining. pp. 215–223. ACM (2017)
7. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Adv. in neural inf. processing systems. pp. 1097–1105 (2012)
8. Lee, J., Magoules, F.: Detection of concept drift for learning from stream data. In: 14th Int. Conf. on High Performance Computing and Communication & 9th Int. Conf. on Embedded Software and Systems. pp. 241–245. IEEE (2012)
9. Lei, P.R.: Mining maritime traffic conflict trajectories from a massive ais data. Knowledge and Inf. Systems pp. 1–27 (2019)
10. Lipton, Z.C.: The mythos of model interpretability. arXiv preprint arXiv:1606.03490 (2016)
11. Lv, J., Li, Q., Sun, Q., Wang, X.: T-conv: a convolutional neural network for multi-scale taxi trajectory prediction. In: 2018 IEEE int. conf. on big data and smart computing (bigcomp). pp. 82–89. IEEE (2018)
12. Manzoor, E., Lamba, H., Akoglu, L.: xstream: Outlier detection in feature-evolving data streams. In: Proc. of the 24th ACM SIGKDD Int. Conf. on Knowledge Discovery & Data Mining. pp. 1963–1972. KDD '18, ACM (2018)
13. Mohammadi, M., Al-Fuqaha, A.: Enabling cognitive smart cities using big data and machine learning: Approaches and challenges. IEEE Communications Magazine **56**(2), 94–101 (2018)
14. Nguyen, D., Vadaine, R., Hajduch, G., Garello, R., Fablet, R.: A multi-task deep learning architecture for maritime surveillance using ais data streams. In: 5th Int. Conf. on Data Science and Advanced Analytics (DSAA). pp. 331–340. IEEE (2018)
15. Nyman, E.: Techno-optimism and ocean governance: New trends in maritime monitoring. Marine Policy **99**, 30–33 (2019)
16. Patroumpas, K., Artikis, A., Katzouris, N., Vodas, M., Theodoridis, Y., Pelekis, N.: Event recognition for maritime surveillance. In: Adv. in Database Technology - EDBT. pp. 629–640 (2015)
17. Qahtan, A.A., Alharbi, B., Wang, S., Zhang, X.: A pca-based change detection framework for multidimensional data streams: Change detection in multidimensional data streams. In: Proc. of the 21th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining. pp. 935–944. ACM (2015)
18. Soares, A., Dividino, R., Abreu, F., Brousseau, M., Isenor, A.W., Webb, S., Matwin, S.: Crisis: Integrating ais and ocean data streams using semantic web standards for event detection. In: Int. Conf. on Military Communications and Inf. Systems (2019a)
19. Wen, Y., Zhang, Y., Huang, L., Zhou, C., Xiao, C., Zhang, F., Peng, X., Zhan, W., Sui, Z.: Semantic modelling of ship behavior in harbor based on ontology and dynamic bayesian network. ISPRS Int. Journal of Geo-Inf. **8**(3), 107 (2019)
20. Zhao, L., Shi, G.: Maritime anomaly detection using density-based clustering and recurrent neural network. The Journal of Navigation **72**(4), 894–916 (2019)

21. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proc. of the IEEE conf. on computer vision and pattern recognition. pp. 2921–2929 (2016)