

Learning Motion Patterns in AIS Data and Detecting Anomalous Vessel Behavior

Anton Kullberg, Isaac Skog and Gustaf Hendeby

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-183212>

N.B.: When citing this work, cite the original publication.

Kullberg, A., Skog, I., Hendeby, G., (2021), Learning Motion Patterns in AIS Data and Detecting Anomalous Vessel Behavior, *2021 IEEE 24th International Conference on Information Fusion (FUSION)*, , 1-8. <https://doi.org/10.23919/FUSION49465.2021.9627027>

Original publication available at:

<https://doi.org/10.23919/FUSION49465.2021.9627027>

Copyright: Institute of Electrical and Electronics Engineers

<http://www.ieee.org/index.html>

©2022 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Learning Motion Patterns in AIS Data and Detecting Anomalous Vessel Behavior

Anton Kullberg*, Isaac Skog*[†], *Senior Member, IEEE*, and Gustaf Hendeby*, *Senior Member, IEEE*.

*Div. Automatic Control, Linköping University, Linköping, Sweden

Email: {anton.kullberg, isaac.skog, gustaf.hendeby}@liu.se

[†]Dept. Underwater Technology FOI Swedish Defence Research Agency, Kista, Sweden

Abstract—A new approach to anomaly detection in maritime traffic based on Automatic Identification System (AIS) data is proposed. The method recursively learns a model of the nominal vessel routes from AIS data and simultaneously estimates the current state of the vessels. A distinction between anomalies and measurement outliers is made and a method to detect and distinguish between the two is proposed. The anomaly and outlier detection is based on statistical testing relative to the current motion model. The proposed method is evaluated on historical AIS data from a coastal area in Sweden and is shown to detect previously unseen motions.

Index Terms—Marine Safety, Anomaly Detection, Outlier Detection

I. INTRODUCTION

The maritime sector currently comprises around 80% of the annual international trade according to the United Nations. Hence, the safety and security of the maritime sector is of utmost importance. Vessel Traffic Monitoring and Information Systems (VTMSIS) are systems designed for active monitoring of maritime traffic. VTMSIS are particularly used for estimating vessel trajectories and detecting anomalous behavior, such as vessels not adhering to established fairways or deviating from an expected trajectory.

Anomaly detection in the maritime domain is of interest for several purposes. Firstly, anomalous behavior can be due to a mechanical fault of a vessel and could therefore require a response from sea rescue. Secondly, anomalous behavior can be an indirect indication of illicit activity, as explored in [1]. Particularly, vessels might take an unexpected path to avoid pirating or because their is some sort of blockade.

To be able to detect and distinguish anomalous behavior, it is essential to establish an accurate model of how vessels typically behaves. One way of establishing traffic patterns was investigated in [2], where historical AIS data was used to construct “routes” in the form of a graph, which were then used for trajectory prediction and anomaly detection. Routes were dynamically added to the model over time. However, the inclusion of a new route was heuristically determined if a fixed number of points did not belong to any of the previously known routes.

Under the assumption of an accurate model, either physically derived or learned from data, the predictions that the

model produces can be used to determine whether a system, i.e., a vessel, is behaving “correctly” or not. The process of detecting anomalies in a model is highly related to outlier rejection, and the two terms are often used interchangeably in the literature. However, in this paper a distinction is made between anomalies and outliers. Outliers are due to errors in the sensor readings or disturbances in the input to the sensor. Anomalies are the result of the system not following the motion model. From a model perspective, the latter can be seen as a change in motion dynamics, which will create a mismatch between the model predictions and the sensor readings (observations). However, contrary to outliers that typically cause sporadic deviations between the observations and the model predictions, the anomalies typically cause long-term deviations. It is stressed that anomalous behavior is only really anomalous if it is not repeated. If the anomalous behavior occurs several times, it should be deemed normal or typical. Hence, it is essential to adapt the model over time to capture the current typical behavior. Furthermore, it is stressed that the interest here is to detect anomalous motion, rather than specific faults (anomalies) related to AIS data, such as on-off switching or power outages, as was treated in [3]. Such anomalies are related to deviations in the sampling interval, which is not treated here.

The traffic monitoring problem essentially consists of two parts: state estimation and detection of anomalous motion. Standard approaches to state estimation is to use Bayesian filters, such as the Kalman filter or particle filters. Generally, these approaches specify a state transition model, which describes the system dynamics. Such a model is often constructed based on physical insight or through the use of system identification techniques [4]. Adaptive models that can incorporate new information online, are particularly interesting in the context of this paper. In this way, typical vessel behavior can be learned recursively and the model can adapt to possible changes in the nominal behavior, which can be naturally incorporated in the model. One such type of model is the Gaussian process state-space model with either partly, or completely, unknown dynamics, see e.g. [5].

In [6], we presented a model consisting of an *a priori* assumed simple, motion model augmented with a basis function expansion, representing unknown motion dynamics. By allowing the unknown dynamics to be spatially dependent, typical motion patterns are learned online. In this paper, the

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

framework in [6] is used to develop a new approach to anomalous behavior detection in maritime traffic based on AIS data. The method learns vessel routes recursively and thus improves over time, allowing new routes to be adopted naturally within the framework. Furthermore, anomaly detection is integrated into the same framework and particularly distinguishes between outliers (single erroneous measurements) and anomalies (vessels not following the learned motion model).

II. PROBLEM FORMULATION

The traffic monitoring problem consists of two parts: state estimation and anomaly detection, which are here treated separately. For both state estimation and anomaly detection, AIS data is used. AIS messages consists of dynamic information, such as position, as well as static information, such as a Maritime Mobile Service Identity (MMSI) number. For the purposes here, the position, speed, heading, time, and vessel MMSI are used. The MMSI is used for identification purposes and acts as a label for the measurements to avoid the data association problem, which was treated in [7].

Concerns have been raised about using the MMSI number as an absolute identification for AIS messages because of the possibility for ships to change MMSI. However, changing one's MMSI is often associated with illicit activity, which is assumed to be non-existent in the surveilled area (the bay area is essentially only trafficked by large freighting/passenger ships as well as small hobby vessels). As such, the MMSI number should serve as a viable identifier for the context of this paper. Moreover, there is also the concern regarding spoofing of the spatial content of an AIS message to appear to be somewhere else, see [8]. Using the same reasoning as previously, this is assumed to be non-existent in the data. However, the algorithm presented herein should be capable of treating such spoofing, even though it is not treated here. Particularly, such behavior would be detected as an anomaly if the spoofed trajectory is not coherent with the learned model.

A. Outlier and Anomaly Detection

Before proceeding, a few assumptions are necessary to ascertain what anomalous behavior actually means. Firstly, it is assumed that the underlying data has been generated by a "true" model and that the true model is contained in the model set. Secondly, the data is assumed to be rich enough such that the estimated model parameters θ asymptotically approach the true parameters, i.e., assuming that the true model parameters are given by θ^0 , the estimated parameters $\theta_N \rightarrow \theta^0$, as $N \rightarrow \infty$, where N is the number of data points. This assumption does not make any practical difference, but it aids in interpreting the results, as the anomalous behavior is defined in relation to the model used, rather than the data itself.

As alluded to in the introduction, outliers and anomalies are treated as two separate things within the context of this paper. They are both, in a sense, related to a modeling error. Outliers are an error of the sensor readings whereas anomalies relate to motion model errors. They can both be detected in

similar ways, using statistical hypothesis tests. One commonly used test, is the T-score or χ^2 statistic. If the observation \mathbf{y} is assumed normally distributed, then

$$T_k = (\mathbf{y}_k - \hat{\mathbf{y}}_k)^T \mathbf{S}_k^{-1} (\mathbf{y}_k - \hat{\mathbf{y}}_k) \sim \chi^2(n_{\mathbf{y}}), \quad (1)$$

where \mathbf{S}_k denotes the innovation covariance, $\hat{\mathbf{y}}$ is the unbiased predicted measurement, and $n_{\mathbf{y}} = \dim(\mathbf{y})$. By using the inverse Cumulative Distribution Function (CDF) of the χ^2 distribution, a threshold, β , can be chosen for some confidence level α such that $P(T_k < \beta) = \alpha$.

Assuming that the motion model is correct, (1) is a reasonable test statistic for outlier detection. However, if this assumption is relaxed, (1) can no longer be used to distinguish whether the error is caused by an incorrect motion model or observation model. Still, note that an outlier is typically an instantaneous phenomenon, whereas anomalous behavior normally occurs over a longer time span. Essentially, during anomalous system behavior, the state transition dynamics will induce temporal correlations in the predictive residuals and there will be several consecutive deviations between prediction $\hat{\mathbf{y}}_k$ and observation \mathbf{y}_k . Hence, one way of detecting an anomaly is via the mean of several consecutive T-scores, i.e.,

$$T_{k-\beta+1:k} = \frac{1}{\beta} \sum_{i=k-\beta+1}^k T_i, \quad (2)$$

where β is the window length. Since T_i and T_j , $i \neq j$, are independent, then

$$\beta T_{k-\beta+1:k} \sim \chi^2(\beta n_{\mathbf{y}}). \quad (3)$$

This can also be implemented recursively with an exponential forgetting factor as

$$T_\lambda = \lambda T_\lambda + (1 - \lambda) T_k, \quad (4)$$

where λ is the forgetting factor, chosen to correspond to some particular window length β . Note that T_λ is a running mean and is approximately $\chi^2(n_{\mathbf{y}})$ distributed.

Another way of detecting anomalous behavior is through the use of n -step prediction. Assuming a correct motion model, simulating the model n steps should still produce a valid observation estimate $\hat{\mathbf{y}}_k$. Hence, using n -step prediction with either (1) or (4) can also be used to detect anomalies. Note that n -step prediction combined with (1) is also sensitive to outliers and cannot discern between the two. All the aforementioned techniques are investigated and evaluated within this paper.

B. State-Space Model

State estimation is dealt with in the form of a state-space model augmented with a basis function expansion. This model is then used in an Extended Kalman Filter (EKF) to estimate the states of the vessels as well as infer typical motion patterns, which is captured by the basis function expansion. A general state-space model augmented with an unknown function is given by

$$\mathbf{x}_{k+1} = \mathbf{f}_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k, \mathbf{g}_k) \quad (5a)$$

$$\mathbf{y}_k = \mathbf{h}_k(\mathbf{x}_k, \mathbf{e}_k). \quad (5b)$$

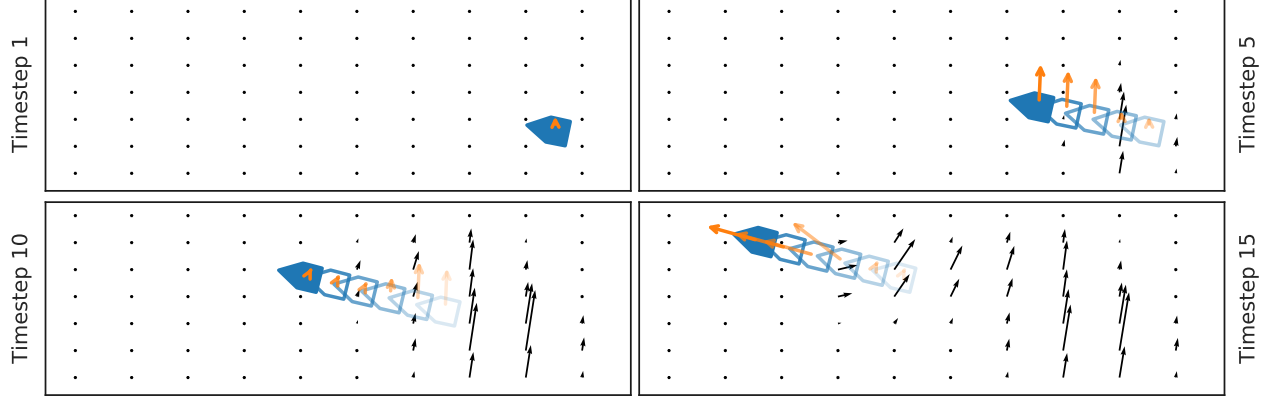


Fig. 1: Recursive learning of the ship motion patterns. The ship is shown in blue, the learned function \mathbf{g}_k is shown in black and the orange arrows are Cartesian accelerations from a numerical difference scheme at each time step, which is specified on each separate y-axis. Historical positions and accelerations are opaque. Note that the arrows are not equally scaled.

Here, \mathbf{x}_k , \mathbf{u}_k , and \mathbf{y}_k denote the system state, input, and measurement, respectively. Further, \mathbf{f}_k and \mathbf{h}_k are the state transition and measurement model, respectively. Moreover, \mathbf{w}_k and \mathbf{e}_k denote the process and measurement noise, respectively. They are assumed to be mutually uncorrelated white noise processes with covariance matrices \mathbf{Q} and \mathbf{R} , respectively. Lastly, \mathbf{g}_k denotes an unknown input, acting on the system, which is to be inferred from data. Particularly, the unknown input is assumed to be a function of the state of the system, *i.e.*, $\mathbf{g}_k = \mathbf{g}_k(\mathbf{x}_k)$. In this way, the motion dynamics of the vessels, apparent from the historical data, can be incorporated by letting \mathbf{g}_k depend on the position of the system. Particularly, it enables a filtering approach that recursively learns the motion patterns in the data. Further, with this formulation, the model naturally fits into other contexts, such as motion planning for collision avoidance where the model can be used for predicting future trajectories of other ships nearby. Collision avoidance with the aid of AIS data was treated in, *e.g.* [9].

The state transition model (5a) consists of an *a priori* specified motion model, as well as an unknown input \mathbf{g}_k . For the purposes of this paper, a simple, motion model is selected, as fine grained motion dynamics such as sway and roll are not of interest. The state of the model is parametrized with $\mathbf{x} = [p_x \ p_y \ v \ \rho]^\top$. The heading ρ and velocity v are assumed constant and the positional coordinates p_x and p_y are in a local Universal Transverse Mercator (UTM) coordinate system. The prior motion model in discrete form is then given by [10]

$$\mathbf{x}_{k+1} = \underbrace{\begin{bmatrix} 1 & 0 & T_s \cos(\rho) & 0 \\ 0 & 1 & T_s \sin(\rho) & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\mathbf{F}(\mathbf{x}_k)} \mathbf{x}_k + \underbrace{\begin{bmatrix} \cos(\rho)T_s^2/2 & 0 \\ \sin(\rho)T_s^2/2 & 0 \\ T_s & 0 \\ 0 & T_s \end{bmatrix}}_{\mathbf{G}(\mathbf{x}_k)} \tilde{\mathbf{w}}_k, \quad (6)$$

where T_s is the sampling interval. Here, the process noise $\tilde{\mathbf{w}}_k$ contains both deterministic components, relating to parts of the system that have not been explicitly modeled, as well as stochastic components, relating to unknown disturbances. Essentially, $\tilde{\mathbf{w}}_k$ can be described as

$$\tilde{\mathbf{w}}_k = \mathbf{g}_k + \mathbf{w}_k, \quad (7)$$

where \mathbf{g}_k is a deterministic function and \mathbf{w}_k is a white noise process with covariance matrix \mathbf{Q} . This enables an explicit interpretation of what the model is learning. In Fig. 1, a vessel is traveling in a region where, *a priori*, nothing is known about the motion patterns. The “true” Cartesian acceleration of the vessel is visualized with orange arrows, which represents the deterministic parts of $\tilde{\mathbf{w}}_k$, *i.e.*, it is what \mathbf{g}_k models. In black, the learned \mathbf{g}_k is visualized, which captures the acceleration of the vessel. The next time a vessel travels through the region, much is already known about how the vessel should accelerate. Hence, the orange arrows will essentially become smaller and smaller until the deterministic components have been absorbed by \mathbf{g}_k .

In (6), this is a particularly useful representation as \mathbf{g}_k can then be interpreted as the vessel aligned acceleration and angular rate, whereas \mathbf{w}_k only captures unknown/random disturbances acting on the system. The complete motion model is then

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k)\mathbf{x}_k + \mathbf{G}(\mathbf{x}_k)(\mathbf{g}_k + \mathbf{w}_k). \quad (8)$$

C. Basis Function Expansion

To be able to infer the unknown function \mathbf{g}_k online, it is modeled as a basis function expansion, *i.e.*,

$$[\mathbf{g}_k]_j = (\boldsymbol{\phi}_k^j)^\top \boldsymbol{\theta}_k^j \quad j = 1, \dots, J \quad (9)$$

where $[\mathbf{g}_k]_j$ denotes the j :th component of \mathbf{g}_k . Further,

$$\begin{aligned}\phi_k^j &= \left(\phi_1^j \quad \dots \quad \phi_{n_\theta^j}^j \right)^\top, \\ \theta_k^j &= \left(\theta_{k,1}^j \quad \dots \quad \theta_{k,n_\theta^j}^j \right)^\top,\end{aligned}$$

where $\phi_i^j = \phi_i^j(\mathbf{x}_k)$ is the i :th basis function corresponding to component j of \mathbf{g}_k and $\theta_{k,i}^j$ is the weight of this basis function at time k . Moreover, the dependence of \mathbf{g}_k on \mathbf{x}_k is solely through the basis functions ϕ_i^j .

The basis function can be chosen in a variety of ways. For computational reasons highlighted in [11], the Wendland function [12]

$$\phi_i^j(r_{k,i}^j) = (1 - r_{k,i}^j)_+^6 (35(r_{k,i}^j)^2 + 18r_{k,i}^j + 3)/3, \quad (10)$$

is used, where $(\cdot)_+ = \max(0, \cdot)$ and $r_{k,i}^j = \|\mathbf{z}_k - \xi_i^j\|$ and ξ_i^j is the ‘‘center’’ of basis function i for component j of \mathbf{g}_k . The function (10) is an example of a Compact Support Radial Basis Function (CSRBF), which lends itself to computationally efficient implementation of the joint state estimation and learning procedure, see [11]. Particularly, the support of the function is the unit sphere, which can be altered by multiplying $r_{k,i}^j$ by a factor α .

With the basis function (10), the basis function expansion can be rewritten as

$$[\mathbf{g}_k]_j = (\mathbf{A}_k^j \phi_k^j)^\top \mathbf{A}_k^j \theta_k^j = (\phi_k^j)^\top (\mathbf{A}_k^j)^\top \mathbf{A}_k^j \theta_k^j, \quad (11)$$

where

$$\mathbf{A}_k^j(\mathbf{x}_k) \in \mathbb{R}^{\tilde{n}_\theta^j \times n_\theta^j}$$

is a matrix-valued function that takes \mathbf{x}_k as input and produces an indicator matrix such that only the *active*, or non-zero, basis functions and weights remain.

Lastly, it is assumed that the basis functions are the same for all of the components of \mathbf{g}_k , *i.e.*, $\phi_k^1 = \dots = \phi_k^J$, and as such, the basis function expansion can finally be written as

$$\mathbf{g}_k \triangleq \Phi_k \boldsymbol{\theta}_k = (\mathbf{I} \otimes \phi_k^\top(\mathbf{x}_k) \mathbf{A}_k^\top) \mathbf{A}_k \boldsymbol{\theta}_k \quad (12)$$

To facilitate online learning of \mathbf{g}_k , the weights $\boldsymbol{\theta}_k$ are assumed to follow a random walk described by the covariance matrix $\boldsymbol{\Sigma}_k$ and an augmented state vector $\mathbf{x}_k^e = [\mathbf{x}_k^\top \quad \boldsymbol{\theta}_k^\top]^\top$ is constructed.

D. Complete Model

The AIS data provides measurements of the position, velocity, and heading of the vessels. The heading is transformed to a trigonometric representation to avoid ambiguities around $-\pi$ and π . With the particular state parametrization described by (6), this results in the measurement model

$$\mathbf{y}_k = \underbrace{\begin{bmatrix} [\mathbf{x}_k]_0 & [\mathbf{x}_k]_1 & [\mathbf{x}_k]_2 & \sin[\mathbf{x}_k]_3 & \cos[\mathbf{x}_k]_3 \end{bmatrix}^\top}_{\mathbf{h}(\mathbf{x}_k)} + \mathbf{e}_k, \quad (13)$$

where $[\mathbf{x}_k]_i$ refers to the i :th component of \mathbf{x}_k .

Lastly, to facilitate dependence of \mathbf{g}_k on but a few of the states, \mathbf{x}_k is transformed before entering \mathbf{g}_k . With slight abuse of notation, \mathbf{g}_k is thus defined as

$$\mathbf{g}_k = \mathbf{g}_k(\mathbf{D}\mathbf{x}_k), \quad (14)$$

where \mathbf{D} is constructed to extract the position and heading from \mathbf{x}_k , for reasons which shall become clear later.

The complete model, combining the prior motion model, basis function augmentation, and measurement model, is then given by

$$\mathbf{x}_{k+1}^e \triangleq \mathbf{f}_k^e(\mathbf{x}_k^e, \mathbf{u}_k, \mathbf{w}_k^e, \mathbf{g}_k) \quad (15a)$$

$$\mathbf{f}_k^e \triangleq \begin{bmatrix} \mathbf{F}(\mathbf{x}_k) \mathbf{x}_k + \mathbf{G}(\mathbf{x}_k) \mathbf{g}_k + \tilde{\mathbf{w}}_k \\ \boldsymbol{\theta}_k + \hat{\mathbf{w}}_k \end{bmatrix} \quad (15b)$$

$$\mathbf{g}_k = (\mathbf{I} \otimes \phi_k^\top(\mathbf{D}\mathbf{x}_k) \mathbf{A}_k^\top) \mathbf{A}_k \boldsymbol{\theta}_k \triangleq \Phi_k \boldsymbol{\theta}_k \quad (15c)$$

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{e}_k, \quad (15d)$$

where \mathbf{e}_k , $\tilde{\mathbf{w}}_k$ and $\hat{\mathbf{w}}_k$ are mutually independent white noise processes with covariance matrices \mathbf{R}_k , \mathbf{Q}_k , and $\boldsymbol{\Sigma}_k$, respectively.

III. ESTIMATION

To estimate the states of the vessels and learn the parameters of \mathbf{g}_k , an EKF is used. The EKF propagates the first two moments, *i.e.*, the mean and the covariance of the states and the basis function weights. The algorithm is summarized in Algorithm 1, where superscript MMSI indicates what MMSI the state estimate and state error covariance belongs to. The necessary matrices \mathbf{F}_k , and \mathbf{G}_k , and the vector $\mathbf{h}(\mathbf{x}_k)$ are given by (6) and (13), respectively. The state and covariance of ships not seen before are initialized as $\hat{\mathbf{x}}_0$, \mathbf{P}_0^{xx} , and $\mathbf{P}_0^{x\theta}$. This prior is centered on the first measurement with an initial uncertainty given by $\mathbf{P}_0^{xx} = \mathbf{R}_k$ and $\mathbf{P}_0^{x\theta} = \mathbf{0}$. For more details, see [11].

IV. SCENARIO DESCRIPTION

The considered surveillance region is a coastal area outside of Västervik, south-east Sweden, see Fig. 2. The region was selected because it has a port and several fairways. Hence, both anomalous behavior detection as well as trajectory prediction is of interest.

A. AIS Data

The data consists of AIS messages from traffic in the region, between January 2019 and April 2020. The complete data set consists of 175 485 messages. The observed traffic data originates from a variety of ship types. As suggested by [2], for learning the nominal travel patterns, only the data from ship types, such as large freighters, fishing, and passenger ships, that under normal operation follows established fairways should be used. The ship types considered in this paper are summarized in Table I. For more details on ship types, see [13].

Furthermore, the rate of AIS-messages typically varies depending on, for instance, the vessel type and current speed. For purely practical reasons and to simplify the resulting algorithm, the AIS data was resampled to regular sampling

intervals for each vessel in the data. The resampling procedure can be summarized as follows:

- 1) Group the data w.r.t. ship callsign.
- 2) Loop over each group and:
 - a) Calculate the sampling interval as time difference between consecutive samples.
 - b) Calculate difference in sampling interval between consecutive samples.
 - c) Split the group into subgroups if the sampling interval difference is “large”.
 - d) Resample each subgroup to the minimum sampling interval within the subgroup or at most $T_s = 20$ s.

Essentially, the data then consists of groups where each group is regularly sampled. Note that this does not mean that each vessel has a fixed sampling interval. It does, however, render “consecutive” measurements to be equally spaced, where “consecutive” depends on the particular meaning of “large”. The sampling interval restriction is to ensure “activation” of the basis functions and to be able to use the learned function \mathbf{g}_k effectively. Lastly, measurements with zero speed are removed as they are deemed uninteresting for the purposes herein. The resulting data consists of 9 058 messages.

B. Choice of Basis Function Centers

The basis functions need to be placed on a grid in the region. The motion patterns of the ships are expected to differ depending on if a ship is heading in or out of port. Therefore, the basis function centers need to be placed both spatially and with regards to heading. Hence, the centers $\xi \in \mathbb{R}^3$, as the function \mathbf{g}_k must depend on both position and heading. Spatially, they are placed as seen in Fig. 2, empirically determined from the minimum and maximum of the latitudinal and longitudinal data. The centers are spatially separated by 25 meters, which was empirically determined to be sufficient, given the slow dynamics of the large ships. For the heading, 6 basis functions are used, equally spaced on the unit circle. To make efficient use of the heading information, a trigonometric encoding is used. This representation is encoded in the basis functions themselves, as in [14]. Particularly, the radius calculation is altered to

$$r_{k,i}^j = \|\gamma(\mathbf{x}_k) - \gamma(\xi_i^j)\|, \quad (16)$$

where $\gamma(\mathbf{x}) = [[\mathbf{x}]_0 \quad [\mathbf{x}]_1 \quad \sin([\mathbf{x}]_3) \quad \cos([\mathbf{x}]_3)]^\top$. Hence, a trigonometric representation of the heading is used whereas the positional information remains the same. It is important to realize that the basis function centers themselves are encoded in the heading domain, whereas the evaluation occurs in the higher dimensional space described by the γ projection, so that the three-dimensional basis function grid is retained.

V. RESULTS

The algorithm was first evaluated on the entire data set to learn the common behavioral patterns of the ships. The process noise covariance is set to $\mathbf{Q} = \text{diag}[0.01 \text{ (m/s}^2\text{)}^2 \ 0.1 \text{ (rad/s)}^2]$ and the measurement noise

Algorithm 1: Joint state inference and model learning[11]

Result: $\hat{\mathbf{x}}_{N|N}^{\text{MMSI}}, \hat{\boldsymbol{\theta}}_{N|N}, \mathbf{P}^{xx, \text{MMSI}}_{N|N}, \mathbf{P}^{x\theta, \text{MMSI}}_{N|N}, \mathbf{P}^{\theta\theta}_{N|N}$
Input: $\mathbf{y}_{1:N}, \hat{\mathbf{x}}_0, \mathbf{P}_0^{xx}, \mathbf{P}_0^{x\theta}$
for $k=1:N$ **do**
 MMSI = $\mathbf{y}_k^{\text{MMSI}}$ // Get MMSI from observation
 if $\hat{\mathbf{x}}_{k-1|k-1}^{\text{MMSI}}$ *exists* **then**
 $\hat{\mathbf{x}}_{k-1|k-1} := \hat{\mathbf{x}}_{k-1|k-1}^{\text{MMSI}}$
 $\mathbf{P}_{k-1|k-1}^{xx} := \mathbf{P}_{k-1|k-1}^{xx, \text{MMSI}}$
 $\mathbf{P}_{k-1|k-1}^{x\theta} := \mathbf{P}_{k-1|k-1}^{x\theta, \text{MMSI}}$
 else
 $\hat{\mathbf{x}}_{k-1|k-1} := \hat{\mathbf{x}}_0$
 $\mathbf{P}_{k-1|k-1}^{xx} := \mathbf{P}_0^{xx}$
 $\mathbf{P}_{k-1|k-1}^{x\theta} := \mathbf{P}_0^{x\theta}$
 end
 // Propagate forward in time
 $\mathbf{P}_{k-1|k-1} = \begin{bmatrix} \mathbf{P}_{k-1|k-1}^{xx} & \mathbf{P}_{k-1|k-1}^{x\theta} \\ \mathbf{P}_{k-1|k-1}^{x\theta} & \mathbf{P}_{k-1|k-1}^{\theta\theta} \end{bmatrix}$
 $\hat{\mathbf{g}}_{k-1} = \Phi_k(\hat{\mathbf{x}}_{k-1|k-1})\hat{\boldsymbol{\theta}}_{k-1|k-1}$
 $\hat{\mathbf{x}}_{k|k-1} = \mathbf{F}_k\hat{\mathbf{x}}_{k-1|k-1} + \mathbf{G}_k\hat{\mathbf{g}}_{k-1}$
 $\hat{\boldsymbol{\theta}}_{k|k-1} = \hat{\boldsymbol{\theta}}_{k-1|k-1}$
 $\mathbf{P}_{k|k-1} = \nabla_{\mathbf{x}^e}\mathbf{F}_k\mathbf{P}_{k|k-1}(\nabla_{\mathbf{x}^e}\mathbf{F}_k)^\top + \mathbf{Q}_k^e$
 // Update estimate of current MMSI
 $\mathbf{S}_k := \mathbf{R}_k + \mathbf{H}\mathbf{P}_{k|k-1}^{xx}\mathbf{H}^\top$
 $\mathbf{K}_k^x := \mathbf{P}_{k|k-1}^{xx}\mathbf{H}^\top\mathbf{S}_k^{-1}$
 $\mathbf{K}_k^\theta := \mathbf{P}_{k|k-1}^{\theta\theta}\mathbf{H}^\top\mathbf{S}_k^{-1}$
 $\mathbf{K}_k := \begin{bmatrix} \mathbf{K}_k^x & \mathbf{K}_k^\theta \end{bmatrix}^\top$
 $\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k^x(\mathbf{y}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1})$
 $\hat{\boldsymbol{\theta}}_{k|k} = \hat{\boldsymbol{\theta}}_{k|k-1} + \mathbf{A}_k^\top\mathbf{A}_k\mathbf{K}_k^\theta(\mathbf{y}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1})$
 $\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H})\mathbf{P}_{k|k-1}(\mathbf{I} - \mathbf{K}_k\mathbf{H})^\top + \mathbf{K}_k\mathbf{R}_k\mathbf{K}_k^\top$
 // Save updated estimate
 $\hat{\mathbf{x}}_{k|k}^{\text{MMSI}} := \hat{\mathbf{x}}_{k|k}$
 $\mathbf{P}_{k|k}^{xx, \text{MMSI}} := \mathbf{P}_{k|k}^{xx}$
 $\mathbf{P}_{k|k}^{x\theta, \text{MMSI}} := \mathbf{P}_{k|k}^{x\theta}$
end

TABLE I
AIS SHIP TYPES

Ship type	Type name
30	Fishing
60–69	Passenger
70–79	Cargo
80–89	Tanker
90	Other

covariance is $\mathbf{R} = \text{diag}[5 \text{ m}^2, 5 \text{ m}^2, 0.5 \text{ (m/s)}^2, 0.1, 0.1]$. The prior uncertainty is $\mathbf{P}_0^{\theta\theta} = 0.001\mathbf{I}$ for the basis function weights and $\mathbf{P}_0^{xx} = 10\mathbf{I}$ for the vessel states. The basis function scaling is $2.5\delta_c$ for the positional components and 0.75 for the angular components. The exponential forgetting




Fig. 2: The coastal region outside of Västervik, Sweden, where the AIS data was collected. Image gathered from Google Maps, [15]. In orange, the positional basis function centers are visualized. The basis function grid has a spacing of 25 meters.

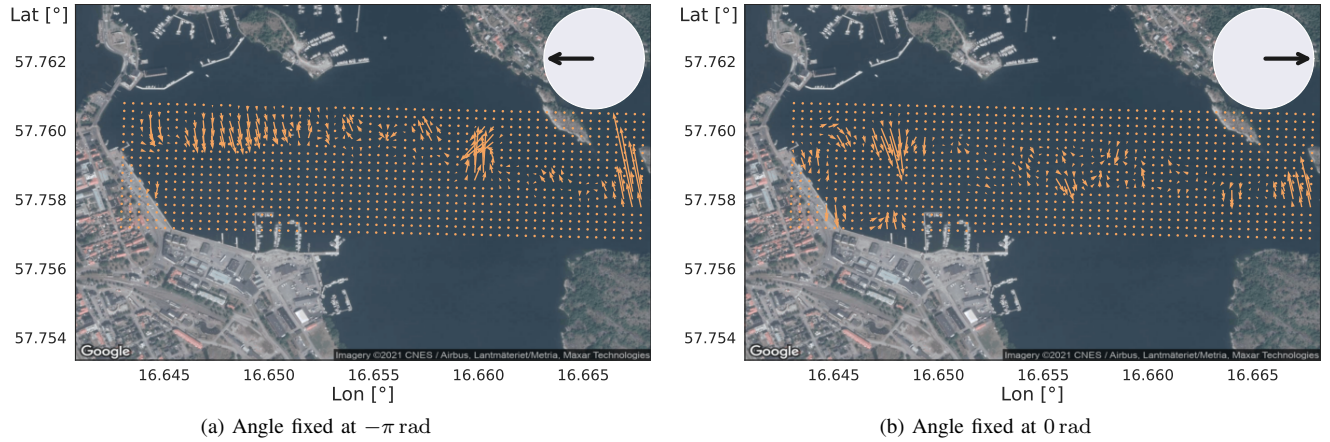


Fig. 3: The learned motion patterns evaluated on a spatial grid for two different angles. Angle is visualized in the top right corner. Note that the arrows are scaled to the image and does not represent absolute accelerations, but can be interpreted relative to each other.

factor $\lambda = 0.5$ which corresponds to a window length of $\beta \approx 3$.

This results in a learned function \mathbf{g}_k which has to be evaluated at specific points to be interpretable. Hence, the heading is fixed and the function is then evaluated on the centers of the basis functions. Further, recall that \mathbf{g}_k represents ship-aligned acceleration and angular velocity. Thus, for interpretability, a projection to Cartesian acceleration is made by the relationship

$$\begin{bmatrix} \dot{v}_x \\ \dot{v}_y \end{bmatrix} = \begin{bmatrix} \cos(\rho) & -v \sin(\rho) \\ \sin(\rho) & v \cos(\rho) \end{bmatrix} \begin{bmatrix} \dot{v} \\ \dot{\rho} \end{bmatrix}. \quad (17)$$

The projection is velocity dependent and as such, the velocity is fixed to the average velocity of the data.

Two such \mathbf{g}_k evaluations are shown in Fig. 3, where the angle is $-\pi$ rad for Fig. 3(a) and 0 rad for Fig. 3(b). Because of the heading dependence of \mathbf{g} , the learned model captures and separates between the behavior coming in to the harbor and going out of the harbor, respectively. In a way, it captures “waypoints” where the vessels usually change

direction, without explicit knowledge of the waypoints. In this sense, the shipways are implicitly captured by \mathbf{g} .

To evaluate the anomaly detection, two anomalies were constructed where the anomalous vessel initially adhered to the typical motion pattern, but then abruptly turned away from the fairway. To evaluate the anomaly detection, the T-score (1) and the forgetting factor batched T-score (2) with both one-step ahead, as well as three-step ahead, prediction was computed for each time step of the anomalous trajectory. As the sampling time varies depending on vessel, the one-step prediction horizons vary between 3-20 seconds and three-step between 9-60 seconds. For reference, the same scores were also computed for vessels adhering to the behavioral patterns. The resulting statistics can be found in Fig. 5. The first and last 10 samples of each vessel trajectory are removed due to transients and boundary behavior of \mathbf{g}_k , respectively. The anomalous trajectories have a clear increase in all of the statistics, whereas the statistics from the vessels moving according to the nominal dynamical model are mostly very




Fig. 4: Sharp loop present only once in the data is detected as an anomaly the second time it is seen.

low and show few spikes. There is one specific occurrence of interest in the metrics corresponding to the nominally correct vessels; around time step 730, callsign “2” experiences a spike in all metrics. This is due to the vessel in question performing an extremely sharp loop, see Fig. 4, which has only been seen once before by the algorithm and it is hence, still an anomaly. Note that even though the T-score (1) appears similar to the forgetting factor T-score (2), it is a point-statistic and can thus not distinguish between an outlier and an anomaly.

VI. CONCLUSION AND FUTURE WORK

An new approach to anomaly detection in maritime traffic based on AIS data has been presented. The method uses a joint state estimation and model learning scheme to learn vessel motion patterns recursively, while maintaining an estimate of the vessel state at all time steps. The nominal behavior is defined relative to the current estimated motion model and anomaly detection is performed by statistical testing with this model. As the model learns recursively, the nominal behavior is updated online and frequently seen behaviors are incorporated naturally within the model. A drawback with the method is that it is memory intensive. On the other hand, historical data can be discarded after being incorporated in the model. Further, the proposed method clearly detects both the artificially produced anomalous behavior, as well as the sharp loop apparent in the original data.

Future work should concern not only anomaly detection, but rejection as well, to accurately select what behaviors should be incorporated in the model.

REFERENCES

- [1] J. R. Watson and A. J. Woodill, “Anticipating illegal maritime activities from anomalous multiscale fleet behaviors,” *arXiv*, no. 80, 2019.
- [2] G. Pallotta, M. Vespe, and K. Bryan, “Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction,” *Entropy*, vol. 15, no. 6, pp. 2218–2245, 2013.
- [3] S. K. Singh and F. Heymann, “On the effectiveness of ai-assisted anomaly detection methods in maritime navigation,” in *FUSION 2020*, Virtual, Jul. 2020.
- [4] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Pearson Education, 1998.
- [5] K. Berntorp, “Online Bayesian Inference and Learning of Gaussian-Process State-Space Models,” *Automatica*, vol. 129, 2021.
- [6] A. Kullberg, I. Skog, and G. Hendeby, “Learning Driver Behaviors Using A Gaussian Process Augmented State-Space Model,” in *FUSION 2020*, Virtual, Jul. 2020.
- [7] J. Y. Yu, “Deep learning approaches for AIS data association in the context of maritime domain awareness,” in *FUSION 2020*, no. July, Virtual, Jul. 2020.
- [8] F. Katsilieris, P. Braca, and S. Coraluppi, “Detection of malicious AIS position spoofing by exploiting radar information,” *Proc. 16th Int. Conf. Inf. Fusion, FUSION 2013*, pp. 1196–1203, 2013.
- [9] E. Tu, G. Zhang, L. Rachmawati, E. Rajabally, and G. B. Huang, “Exploiting AIS Data for Intelligent Maritime Navigation: A Comprehensive Survey from Data to Methodology,” *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 5, pp. 1559–1582, 2018.
- [10] M. Roth, G. Hendeby, and F. Gustafsson, “EKF/UKF Maneuvering Target Tracking using Coordinated Turn Models with Polar Cartesian Velocity,” in *17th Int. Conf. Inf. Fusion*. Salamanca: IEEE, Jul. 2014, pp. 1–8.
- [11] A. Kullberg, I. R. Skog, and G. Hendeby, “Online Joint State Inference and Learning of Partially Unknown State-Space Models,” *IEEE Trans. Signal Process.*, vol. 69, pp. 4149 – 4161, 2021.
- [12] H. Wendland, “On the smoothness of positive definite and radial functions,” *J. Comput. Appl. Math.*, vol. 101, no. 1-2, pp. 177–188, 1999.
- [13] MarineTraffic, “What is the significance of the ais shiptype number,” <https://tinyurl.com/3ucmbwjh>, Accessed Apr. 20, 2021. [Online].
- [14] I. Skog, “Nonintrusive Elevator System Fault Detection Using Learned Traffic Patterns,” *IEEE Sensors Lett.*, vol. 4, no. 11, pp. 11–14, 2020.
- [15] Google, “Västervik, Sweden,” <https://www.google.com/maps>, Accessed Apr. 20, 2021. [Online].

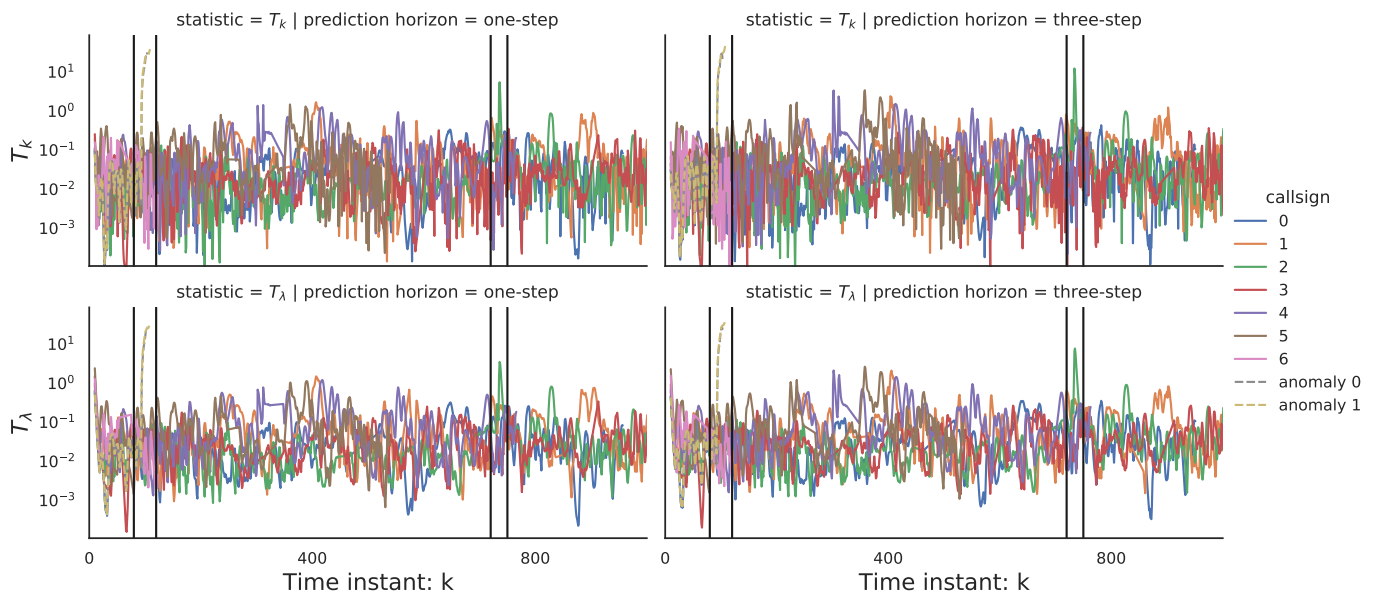


Fig. 5: T-score (1) and forgetting factor T-score (4) statistics for both nominally correct and anomalous vessels when using the learned motion model for anomaly detection. Two regions of interest are marked with vertical black lines. The leftmost is the statistics for the constructed anomalies and the rightmost is the effect of a sharp loop present in the data once, i.e., an anomaly, see Fig. 4.