



HAL
open science

A Multi-task Deep Learning Architecture for Maritime Surveillance using AIS Data Streams

Duong Nguyen, Rodolphe Vadaine, Guillaume Hajduch, René Garello, Ronan Fablet

► **To cite this version:**

Duong Nguyen, Rodolphe Vadaine, Guillaume Hajduch, René Garello, Ronan Fablet. A Multi-task Deep Learning Architecture for Maritime Surveillance using AIS Data Streams. 2018 IEEE International Conference on Data Science and Advanced Analytics (DSAA), Oct 2018, Turin, Italy. 10.1109/DSAA.2018.00044 . hal-01808176v4

HAL Id: hal-01808176

<https://imt-atlantique.hal.science/hal-01808176v4>

Submitted on 15 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Multi-task Deep Learning Architecture for Maritime Surveillance using AIS Data Streams

Duong Nguyen*, Rodolphe Vadaine†, Guillaume Hajduch†, René Garello* and Ronan Fablet*

*IMT Atlantique, Lab-STICC, UBL, 29238 Brest, France

Email: {van.nguyen1, rene.garello, ronan.fablet}@imt-atlantique.fr

†CLS - Space and Ground Segments, 29280 Brest, France

Email: {rvadaine, ghajduch}@cls.fr

Abstract—In a world of global trading, maritime safety, security and efficiency are crucial issues. We propose a multi-task deep learning framework for vessel monitoring using Automatic Identification System (AIS) data streams. We combine recurrent neural networks with latent variable modeling and an embedding of AIS messages to a new representation space to jointly address key issues to be dealt with when considering AIS data streams: massive amount of streaming data, noisy data and irregular time-sampling. We demonstrate the relevance of the proposed deep learning framework on real AIS datasets for a three-task setting, namely trajectory reconstruction, anomaly detection and vessel type identification.

Index Terms—AIS, maritime surveillance, deep learning, transfer learning, trajectory reconstruction, anomaly detection, vessel type identification, variational recurrent neural networks.

I. CONTEXT

In the world of a globalized economy, maritime surveillance is a vital demand. Currently being the most efficient long-distance transporting method, sea transport is carrying about 90% of the world trade¹. With the persistent growth of maritime traffic, safety and security are key issues. Besides, the real-time delivery of maritime situation maps is also necessary for a variety of activities: fishing activities control, smuggling detection, EEZ intrusion detection, transshipment detection, maritime pollution monitoring, etc.

Over the last decades, the development of terrestrial networks and satellite constellations of Automatic Identification System (AIS) has opened a new era in maritime traffic surveillance. Every day, AIS provides on a global scale hundreds of millions of messages [1], which contain ships’ identifier, their Global Positioning System (GPS) coordinates, their speed, course, etc. The potential of this massive amount of data is clearly of interest if tools and models provide means to efficiently extract, detect and analyze relevant information from these data streams. However, current operational systems, which strongly rely on human experts, can only deal with a limited fraction of AIS data.

Thus, the development of AI-based systems is a critical challenge. Beyond the volume of streaming data to be dealt with, there are two other key issues make it difficult to design these types of systems: noise patterns exhibited by AIS data and the irregular time-sampling. Both are very common in AIS

and make the direct application of state-of-the-art supervised machine learning models, including deep learning ones poorly adapted.

This paper addresses these issues and explores deep learning models and architectures, and more specifically Recurrent Neural Networks (RNNs) to develop an automatic system that can process and detect, extract and characterize useful information in AIS data streams for maritime surveillance. More specifically, our key contributions are three-fold:

- The design of a novel big-data-compliant unsupervised architecture which automatically learns and extracts useful information from noisy and partial AIS data streams on a regional scale;
- The joint exploitation of this architecture as a basis for specific tasks using mathematically-sound statistical models, namely trajectory reconstruction and forecasting, maritime route estimation, vessel type identification, detection of abnormal vessel behaviors, etc.;
- The demonstration of the proposed approach’s relevance on real regional datasets off Brittany coast and in the Gulf of Mexico, significantly more complex than case-studies addressed in previous works.

This paper is organized as follows: in Section II, we review the state-of-the-art methods in AIS-based maritime surveillance. Section III provides the background to understand this paper. The proposed method is detailed in Section IV. We present experiments in Section V, and further discuss the main features and performance of our approach in Section VI. Finally, conclusions and perspectives for future work are presented in Section VII.

II. RELATED WORK

In this section, we review the related works in the field of AIS-based maritime traffic surveillance, especially regarding trajectory reconstruction and forecasting and anomaly detection.

Trajectory reconstruction and forecasting: For simplicity purpose, we use here the term “trajectory reconstruction” to refer to both trajectory reconstruction and trajectory forecasting. Early efforts for trajectory reconstruction include linear interpolation, curvilinear interpolation [2] and their improvements [3], [4]. They rely on a physical model of the movement $x_t = x_{t-1} + \Delta t * x'_t$ (where x_t is the position of vessel at

¹According to the International Maritime Organization (IMO)

the time t , x'_t is the deviation of x_t , usually the Speed Over Ground (SOG) and the Course Over Ground (COG)). More sophisticated methods suppose that vessel trajectories follow a distribution and learn it from historical data [5], [6]. Currently, state-of-the-art methods for trajectory reconstruction [7], [8], [9] use the following typical three-step approach: i) the first step involves a clustering method, *e.g.* TRACLUS [10] or TREAD [11] to cluster historical motion data into route patterns, ii) the second one assigns the vessel to be processed to one of these clusters iii) the third one interpolates or predicts the vessel trajectory based on the route pattern of the assigned cluster.

Anomaly detection: Some models detect abnormal behaviors by defining them explicitly [12], [13]. These types of models are usually limited themselves by their own definitions, and can not handle all the complex phenomenons observed at sea. To overcome those drawbacks, other methods detect anomalies implicitly by creating normalcy models, then consider trajectories or trajectory segments that do not suit these models as abnormal. In [14], Rhodes divided the map into small zones and used Normalcy Box to detect abnormal vessel speed in each zone. Gaussian Mixture Models (GMMs), Kernel Density Estimation (KDE) were explored in [15], [16]. More sophisticated methods have used time series analysis techniques, such as Gaussian Process [17], [18], or Bayesian Networks (BNs) [19], [20] to capture the sequential structure of AIS streams. All these models share the same basic idea: in a small region, vessels should perform similar behaviors.

All models and approaches reviewed for trajectory reconstruction and anomaly detection present three main drawbacks:

- They depend on strong priors and can hardly capture all the heterogeneous characteristics of AIS data as well as the varieties of vessels' behaviors. Near-far, fast-slow, etc. are relative definitions and are difficult to be implemented. Almost all current models work only for cargo and tanker vessels on specific high-traffic maritime routes. However, more sophisticated models and relaxed assumptions are required to address the range of vessel types and vessel behaviors revealed by AIS streams on a regional or global scale.
- Most if not all methods exploit at some point a clustering. They typically assume that in specific areas, all vessels tend to perform similar behaviors, and then use clustering methods (Kmeans, DBSCAN, etc.) to find those behaviors. For example, for trajectory reconstruction issues, each cluster is a maritime route [11]; in anomaly detection, each cluster is a speed mode [6], [14], etc. We believe that such clustering steps result in information losses. By contrast, we argue that continuous latent states should be preferred to address the complexity of AIS data streams.
- Current methods do not explicitly address the irregular time-sampling of AIS streams. Non-sequential methods [14] do not take it into account and sequential ones [7] assume they are provided with regularly-sampled streams, which is not true or may result in the creation of

artificial, possibly erroneous AIS positions if interpolation techniques are used as a pre-processing step.

As detailed hereafter, we develop a novel multi-task deep learning framework to address these issues and demonstrate its relevance from experiments on a real AIS dataset on a regional scale.

III. RECURRENT NEURAL NETWORKS WITH LATENT VARIABLES

Deep learning has rapidly become the state-of-the-art framework for a wide range of machine learning problems [21], especially when dealing with large-scale datasets. In this context, recurrent neural networks have showed their advanced abilities for time series processing. Recently, the introduction of recurrent neural networks with latent variables [22], [23], [24] has further extended the application range of RNNs to deal with noisy and heterogeneous sequential data modeling. Latent variables in RNNs provide this type of networks the ability to infer rich state representations and flexible non-linear transition models to account for complex randomness sources [22].

The aim of recurrent neural networks with latent variables is to model the distribution p of a sequence of T observed random variables $x_{t,t=1..T}$. We assume that the generation process of $\{x_t\}$ relies on a sequence of T latent variables $z_{t,t=1..T}$. The joint distribution $p(x_{1:T}, z_{1:T})$ could be factored into:

$$p(x_{1:T}, z_{1:T}) = p_1(x_1, z_1) \prod_{t=2}^T p_t(x_t, z_t | x_{1:t-1}, z_{1:t-1}) \quad (1)$$

$p_t(x_t, z_t | x_{1:t-1}, z_{1:t-1})$ is then factored into two conventional steps of hidden state space models where the prior distribution $p(z_t | x_{1:t-1}, z_{1:t-1})$ denotes the transition function and the conditional distribution $p(x_t | x_{1:t-1}, z_{1:t})$ denotes the emission function. The distribution of the observed data $p(x_{1:T})$ is given by:

$$p(x_{1:T}) = \mathbb{E}_{z_{1:T}} [p(x_{1:T}, z_{1:T})] \quad (2)$$

However, the integral over $z_{1:T}$ in Eq. 2 can not be calculated directly. To circumvent this obstacle, the most common approach [22], [23], [24] is to introduce an approximation $q(z_t | x_{1:t}, z_{1:t-1})$ of the posterior distribution $p(z_t | x_{1:t}, z_{1:t-1})$ then estimate $p(x_{1:T})$ by the Evidence Lower Bound (ELBO):

$$\log p(x_{1:T}) \geq \mathcal{L}(x, p, q) = \mathbb{E}_{z_{1:T} \sim q} [\log p(x_{1:T} | z_{1:T})] - \mathcal{KL}[q(z_{1:T} | x_{1:T}) || p(z_{1:T})] \quad (3)$$

where $\mathcal{KL}[q || p]$ is the Kullback-Leibler divergence between two distributions q and p .

This estimation can be applied for all the family of p and q that can be factored over t :

$$p(x_{1:T} | z_{1:T}) = \prod_{t=1}^T p(x_t | x_{1:t-1}, z_{1:t}) \quad (4)$$

$$q(z_{1:T} | x_{1:T}) = \prod_{t=1}^T q(z_t | x_{1:t}, z_{1:t-1}) \quad (5)$$

$$p(z_{1:T}) = \prod_{t=1}^T p(z_t | x_{1:t-1}, z_{1:t-1}) \quad (6)$$

In this paper, we use Variational Recurrent Neural Networks (VRNNs) of Chung et al. [22]. VRNNs model the historical information $(x_{1:t-1}, z_{1:t-1})$ by the hidden state of its RNN $h_t = h_t(x_{t-1}, z_{t-1}, h_{t-1})$. The conditional distribution $p(x_t | x_{1:t-1}, z_{1:t}) = p(x_t | z_t, h_t)$, the prior distribution $p(z_t | x_{1:t-1}, z_{1:t-1}) = p(z_t | h_t)$ and the variational posterior distribution $q(z_t | x_{1:t}, z_{1:t-1}) = p(z_t | x_t, h_t)$ are parameterized by fully connected networks.

IV. PROPOSED MULTI-TASK VRNN MODEL FOR AIS DATA

As sketched in Fig. 1, we propose a general multi-task Variational Recurrent Neural Network for the analysis of AIS data streams. Received AIS messages are regarded as irregular noisy observations of the true hidden states - called regimes; these regimes themselves may correspond to specific activities (e.g. under way using engine, at anchor, fishing, etc.). The key component of our model is the Embedding block, which converts noisy and irregularly-sampled AIS data to consistent and regularly-sampled hidden regimes. This Embedding block relies on a VRNN [22] and operates at a 10-minute time scale. Higher-level blocks are task-specific submodels, addressing at different time-scales (e.g. daily, monthly,...) the detection of abnormal behaviors, the automatic identification of vessel types, vessel position prediction, the identification of maritime routes, etc.

A. A latent variable model for vessel behaviors

Through a VRNN architecture, we introduce hidden regimes (latents variables) as a data representation² that captures the true maneuvers of vessels (*natural clustering*). Hidden regimes can be regarded as the “roots” of AIS messages. They govern how the vessel moves. From the point of view of higher levels (task-specific layers), hidden regimes provide the necessary information for their task (*hierarchical organization of explanatory factors* and *shared factors across tasks*). They disentangle the underlying information of AIS data (*simplicity of factors dependencies*). For example, saying “this vessel is performing a fishing maneuver” is much more informative than saying “the speed of this vessel is high”.

It is important to note that the hidden regimes are not clusters of AIS messages, because the act of assigning data to group would cause information loss. We share the same idea with [26], that latent variables (hidden regimes in this case) are continuous and there are no simple interpretations of these dimensions.

The introduction of hidden regimes brings us two key benefits: an efficient encoding of AIS datasets and a regularly-sampled sequential representation. Regarding the first aspect, state-of-the-art systems such as TREAD [11] have to store all the AIS messages in the training set, which is updated

²Here we use the criteria defined in [25] to evaluate this representation, readers are encouraged to read [25] before continue.




Fig. 1. Proposed VRNN architecture.

incrementally new AIS messages. Therefore, data volume to be handled for the test phase increases rather linearly with the area of the Region Of Interest (ROI) and the duration of the considered time period. This may prevent such systems from scaling up to regional or global scales. By contrast, once the VRNN is trained, all the knowledge gained from a given AIS dataset is encoded by the characteristics of the hidden regimes, more precisely the fitted conditional distributions $p(z_t | x_{1:t-1}, z_{1:t-1})$ and $p(x_t | x_{1:t-1}, z_{1:t})$. Therefore, for the application of a trained model, there is no need to access the training dataset. This dataset may only be of interest to retrain or fine-tune a given model. It may be noted that the complexity of the representation of the hidden regimes (*i.e.*, the associated number of parameters) does not depend on the training data volume. For instance, in the considered experiments, for a dataset of more than $2 \cdot 10^8$ AIS messages (each message contains several attributes), the fitted hidden regime representation involves about $5 \cdot 10^6$ parameters. The second important feature is the mapping of an input space consisting of a noisy irregularly-sampled time series to a novel regularly-sampled sequential representation which naturally accounts for the different sources of uncertainties exhibited by AIS datasets. Hence, the proposed architecture embeds somehow a time regularization of the input data and does not require the definition of ad hoc denoising and interpolation pre-processing steps, which prove difficult due to the variabilities to be dealt with (e.g., duration of the missing data segments, noise patterns, inhomogeneous space-time variabilities, etc.). From a mathematical point of view, the considered model naturally embeds these issues through the time propagation of the approximate posterior distribution $q(z_t | x_{1:t}, z_{1:t-1})$. Overall, this regularly-sampled sequential representation makes feasible the design of classic architectures on top of the embedding layer to deal with task-specific issues as detailed in Section IV-D, IV-E and IV-F.

B. “Four-hot” representation of AIS messages

Instead of presenting AIS messages directly by their 4-D real-value vector $[lat, lon, SOG, COG]^T$ like methods in the literature [6], we apply the bucketization technique to introduce a novel representation of AIS data: the “four-hot encoding” (Fig. 2). This representation, inspired by the one-hot encoding in language modeling, is created by concatenating the one-hot vectors of 4 attributes in AIS message: latitude




Fig. 2. “Four-hot” vector.

coordinate, longitude coordinate, SOG and COG. To create the one-hot vector of an attribute, we simply divide the entire value range of this attribute into $N_{attribute_i}$ equal-width bins.

The “four-hot encoding” not only brings us the benefits of bucketized presentation but also provides a more structured representation to learn trajectory spatial patterns as illustrated in Section VI. Our four-hot representation shares similarities with [27]. However, in [27], the authors explained their representation as a transformation from feature space to semantic space based on the smoothness prior assumption. They argued that the continuous values of features did not matter, the explanatory factors were the semantic interpretation presented in their discrete vector of these values. We, on the other hand, consider the “four-hot encoding” as a presentation that can i) accelerate the calculation of neural networks (similar to one-hot encoding), ii) disentangle some explanatory factors of input features (see Section VI). The semantic space in our architecture is the space of hidden regimes.

The implicit reduction of the precision of the AIS position and velocity features may be regarded as a drawback of the four-hot representation. We however argue that for the targeted applications there is no need for the embedding block to provide precise numerical features. For example, a speed of 12 knots and a speed of 12.1 knots do not mean any difference in our context.

C. Embedding block

The embedding block is a VRNN [22], where x_t is the “four-hot encoding” of AIS message and z_t is the concatenation of the hidden state of the network and the latent variable at the time t . This layer works at a 10-minute time scale (*i.e.* we downsample AIS data stream to a resolution of 10 minutes) and learns the distribution $p(x_{1:T})$ (via the prior distribution $p(z_t|x_{1:t-1}, z_{1:t-1})$, the generative distribution $p(x_t|x_{1:t-1}, z_{1:t})$ and the approximative posterior distribution $q(z_t|x_{1:t}, z_{1:t-1})$).

After being trained, the embedding layer consistently generates regularly time-sampled hidden regime series. This series is used as input to task-specific submodels as sketched in Fig.1.

D. Trajectory reconstruction submodel

The Embedding block is naturally a generative model, so the construction of a vessel trajectory estimator/predictor on top of this block is relatively direct. We follow the philosophy of [7]. In this approach, one infers maritime contextual information, which is used to enhance the prediction/estimation. The contextual information in [7] was inferred by TREAD [11], which means that each vessel would be assigned to a predefined

route. By contrast, we avoid such a hard assignment to a predefined behavioral cluster. We benefit from the richer contextual representation inferred by the Embedding block. Formally, the proposed trajectory reconstruction model is stated as the inference of the posterior $q(z_t|x_{1:t}, z_{1:t-1})$ and the sampling-resampling from the distribution $p(x_{t+1}|x_{1:t}, z_{1:t}) = \int p(x_{t+1}|x_{1:t}, z_{1:t+1})p(z_{t+1}|x_{1:t}, z_{1:t})dz_{t+1}$ (all learned by the Embedding block) using a particle filter [24].

E. Abnormal behaviour detection submodel

The second specific task on top of the Embedding block is the detection of abnormal behaviors. It comes to define a normalcy model to detect the (unlikely) anomalies w.r.t. this model. As a direct by-product of the trained Embedding block, we can evaluate the likelihood $p(x_{1:T})$ of any input AIS sequence $x_{1:T}$ using a marginalization w.r.t. the hidden regimes. A series of AIS messages with a very low likelihood w.r.t. a given threshold may be regarded as being unlikely for model $p(x_{1:T})$ and hence as abnormal.

One may however consider context-aware detection rules. For example on maritime routes, vessels’ behaviors are roughly identical, which leads to high values for the likelihood $p(x_{1:T})$. In other regions, the variety of vessel types and activities results in much more complex mixtures of behaviors and much lower likelihood values for the normalcy model. The selection of a global threshold over an entire region may not be as appropriate. To address these issues, we introduce a *a contrario* detector [28]³. It works at a 4-hour time scale and addresses the early detection of abnormal vessel behaviors. We divide the map into small cells C_i . In each cell, we calculate the mean m_i and the standard deviation std_i of the $\log p(x_t|x_{1:t-1}, z_{1:t-1})|_{x_t \in C_i}$ using the tracks in the validation set. Any evolution $p(x_t|x_{1:t-1}, z_{1:t-1})$ at timestep t of an AIS track will be considered as an abnormal evolution if its log-likelihood is much lower than the distribution of other log-likelihoods in the same cell. The *a contrario* detection detects if an arbitrary segment is abnormal based on the number of abnormal evolutions in this segment and its length.

F. Vessel type identification submodel

The third task addressed by our model is the identification of the vessel type from its AIS-derived trajectory data. It may be noted that the vessel type should be one of the attributes included in AIS messages. However, not all vessels send their static messages. Some may even send on purpose a false vessel type in AIS messages. A Vessel type identification submodel is then an important tool to detect suspicious behaviors.

Different types of vessels usually perform specific behaviors, which may differ among others in terms of geographical zones, speed patterns, etc. For example, tankers normally follow maritime routes (usually straight lines between two maritime waypoints [11]), their average speed is relatively low, about 12-15 knots, whereas passenger ships have relatively high average speed, about 20-25 knots. If a vessel declares

³We let the reader refer to [28] for a detailed description of the *a contrario* setting.

itself as type “A” but performs a maneuver of type “B”, it is likely that it may carry out illegal activities.

In this study, we design a Vessel type identification sub-model at a 1-day time scale. This submodel explores a Convolutional Neural Network (CNN). The input of this CNN is a $H \times D$ matrix, whose columns are the hidden regimes (dimension H), and D is the number of timesteps in one day (144 in this case). Because the hidden regime is regularly time-sampled, this configuration applies directly.

V. EXPERIMENTS AND RESULTS

We implemented the proposed framework for a three-task model, addressing respectively vessel trajectory reconstruction, abnormal behavior detection and vessel type identification, in the Gulf of Mexico and the abnormal behavior detection off Brittany coast in the Ushant zone⁴. The Ushant water is the entrance to English channel, this region is interesting to maritime surveillance because of its separation scheme and the heavy traffic there. The Gulf of Mexico is relatively large compared to the case-study regions considered in previous studies [11], [15], [17]. This region involves multiple vessel types and activities of vessels. It comprises big ports, fishing zones, oil platforms and dense maritime routes. Overall we considered AIS data from January to March 2017 off Brittany coast in the Ushant zone (2,021,236 AIS messages) and from January to March 2014 in the Gulf of Mexico (180,344,817 AIS messages).

A. Preprocessing

For the pre-processing step, first, infeasible speed or infeasible position messages were removed from the set. To handle the problem of very long sequence when working with RNNs, we split vessels’ tracks into subtracks of from 4 to 24 hours. From now on, in this paper, vessel tracks refer to such subtracks. We also removed tracks whose speed is smaller than 0.1 knots for more than 80% of the time (at anchor or moored vessels).

One objective of the proposed architecture is to deal with irregularly-sampled data. However, we need regularly sampled data to train the model first. In a layman’s term, the Embedding block must see how regularly sampled AIS tracks should be and learn their characteristics, after that (after being trained), it will generate regularly-sampled data from irregularly-sampled ones. Therefore, for the training set, we only chose tracks whose the maximum time interval between two successive received AIS messages is 1 hour, then used constant velocity model to create regularly time-sampled AIS tracks at 10-minute time scale. By doing this, the intervals between two successive AIS messages are small enough that the errors in the estimation of the constant velocity model do not effect our model too much.

⁴The Tensorflow code and the datasets are available at <https://github.com/dnguyengithub/MultitaskAIS>

TABLE I
LOG LIKELIHOODS OF THE EMBEDDING BLOCK WITH DIFFERENT DIMENSION SETTINGS (GULF OF MEXICO DATASET).

Hidden regime dimension	Number of parameters	Log likelihood on training set	Log likelihood on test set
200	1 605 402	-7.592710	-7.678684
400	5 129 202	-6.557936	-7.520255
500	7 611 102	-6.130078	-7.690255

B. Embedding block calibration

We implemented the Embedding block by a VRNN whose the RNN is a single-layer LSTM, the distributions $p(x_t|z_t, h_t)$, $p(z_t|h_t)$, $q(z_t|x_t, h_t)$ are fully connected networks with one hidden layer of the same size of the LSTM’s. $p(x_t|z_t, h_t)$ is binomial, $p(z_t|h_t)$ and $q(z_t|x_t, h_t)$ are Gaussians. The network was trained with stochastic gradient descent using Adam optimizer [29], learning rate of 0.0003.

There is a trade-off between the resolutions of AIS features and the size of the network when choosing the length of the “four-hot encoding”. If the resolutions are too high, the “four-hot” vector will be too long, requires a big hardware memory and computational power; if the resolutions are too low, we lose information. We set here the resolution of the latitude and longitude coordinate at about 1 km, the resolution of SOG at 1 knot and the resolution of COG at 5°. These resolutions are fine enough for almost all the maritime safety, security and efficiency tasks. For example, with this setting, the uncertainty zone of vessel’s position is about 1kmx1km, small enough for position-related tasks.

The choice of the dimension of hidden regime effects the modeling capacity of the Embedding block. As shown in Table I, if the latent size is too small, the model can not capture all the variations of AIS data. In contrast, if the latent size is too big, the model becomes too bulky and overfitting. For the rest of this paper, we set the latent size at 400 for tests on the Gulf of Mexico dataset and at 100 for tests on the Brittany dataset (Ushant water).

C. Vessel trajectory construction

We deleted a 2-hour segment from each AIS track then used the Vessel trajectory construction layer to reconstruct this segment. The maritime contextual information learned by the Embedding block gave the model the ability to reconstruct some complex trajectories like those on the top right and bottom left of Fig. 3. These constructions can not be achieved by interpolation methods such as linear or spline interpolation.

The performance of this layer depends strongly on the maritime contextual information extracted by the Embedding layer. If the extraction is good, the model can predict complicated patterns like those shown in Fig. 3. However, in zones whose the vessel density is low, or in zones where the behaviors of vessels are too complicated for the Embedding layer to learn, the construction layer completely fails to estimate the positions of vessels. In these cases, we use constant velocity method. The switch between particle method




Fig. 3. Trajectory reconstruction examples using the proposed model. Blue dots: received AIS messages; red dots: missing AIS messages; red lines: trajectories reconstructed by our model.

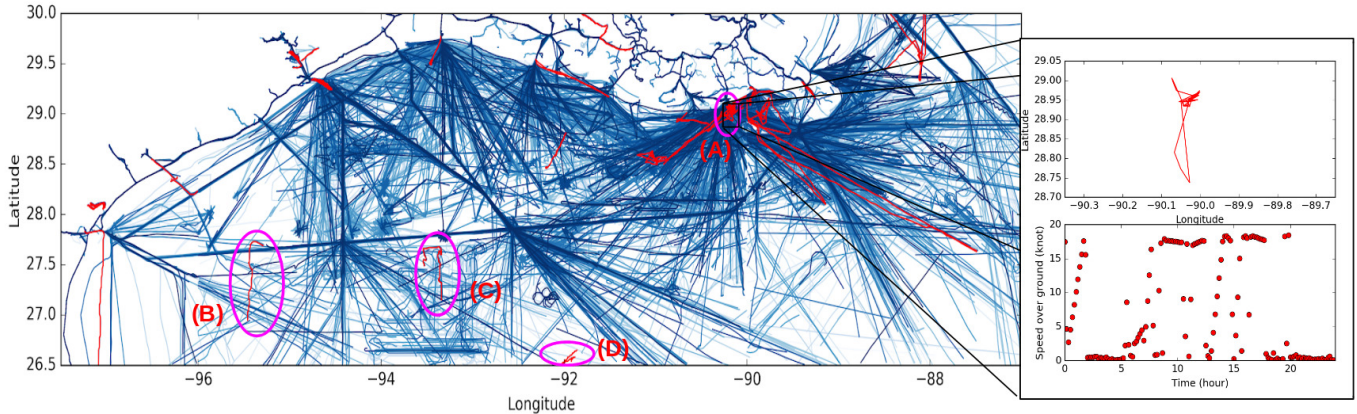


Fig. 4. Detection of abnormal behaviors using global thresholding (Gulf of Mexico dataset). Blue: tracks in the training set (which itself may contain abnormal tracks); red: abnormal tracks detected in the test set. We highlight four examples: a track with an abnormal speed pattern (A) ii), two tracks with abnormal trajectory shapes from others' in the same region (B,C) iii) a track in a low-density area (abnormal zone) (D).

and constant velocity method is automatic, because the model knows when the Embedding layer can not extract the maritime contextual information (based on the value of the probability $p_t(x_t|x_{1:t-1}, z_{1:t-1})$).

D. Abnormal behavior detection

We divided each dataset into 3 sets: a training set to train the model, a validation set to calculate the the *mean* and *std* of the log probability, and a test set to test the anomaly detection. The proportion of the 3 sets was 60/30/10. Although the training sets were used for learning the normalcy model, we did not do data cleaning, *i.e.* the training sets themselves may contain abnormal trajectories. Our framework relies on probabilistic models and implicitly assumes that abnormal trajectories are rare events, that is to say that the probability mass at these trajectories would be very low.

We report the outcome of the anomaly detection submodel when using global threshold detection on the Gulf of Mexico dataset in Fig. 4. A track will be detected as abnormal if its shape is unusual, its speed pattern is rare, or it appears in an abnormal region, etc. Each type of these anomalies corresponds to a signature of trajectory data, like geographical

pattern, geometric pattern, speed and course distribution, etc. These signatures will be presented in Section VI.

For the *a contrario* detection, we split the ROI into small cells of 10kmx10km. The maps of the mean and the standard deviation of the log-likelihood on the Ushant dataset are shown in Fig. 6. We can see that the log-likelihood strongly depends on geographical region, global thresholding would not work. On the mean map, there are some lines/curves of high value, they are the maritime routes. On the maritime routes, the vessel density is high, vessels performs simple and similar maneuvers, so the model can learn these patterns easily. On the other hand, in regions where the vessel density is low, or the behaviors of vessels are too complicated, the identification of abnormal behaviors appears more complex and may require larger training datasets. Detection examples and their corresponding interpretation are shown in Fig. 7.

To evaluate the consistency of the *a contrario* model, we tested this detector for simulated abnormal examples. We translated a normal track out of maritime routes to simulate the divergence from a given route (zone (A)) and translated circle-shaped tracks in zone (B) to zone (C) in Fig. 5 to verify that some specific patterns of vessels' maneuvers should appear in their specific zones. Experiment shows that the model can

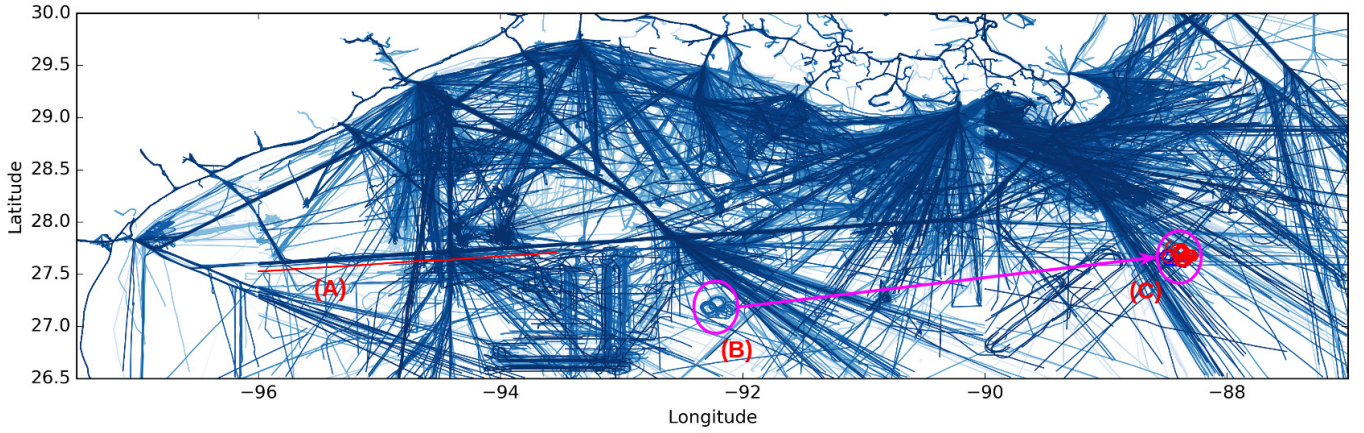



Fig. 5. Example of the *a contrario* anomaly detection on simulated dataset (Gulf of Mexico dataset). The circle-shaped tracks in zones (C) were simulated by translating from (B); (A) is a detection of a divergence from maritime route.



(a) Mean of the log-likelihood in each cell

(b) Std of log-likelihood in each cell

Fig. 6. Maps of the mean and the std of the log-likelihood of the trained model in each cell (Brittany dataset).

detect the divergences if the distance to the maritime route is far enough (10km) and it detects 9 over 13 circle-shaped tracks in zone (C).

In comparison to methods in the literature, our method has several benefits:

- It can detect abnormal patterns that are detected in state-of-the-art methods, such as the double-U-turn detection reported in Fig. 7f and also illustrated in [11].
- Methods like those in [11] and [20] first assign a track to a maritime route, then compare the similarity between this track with the those in the corresponding route to decide whether this track is normal. However, it is very difficult to link tracks like the one in Fig. 7d to a maritime route. Therefore, our model which does not require the prior identification of maritime routes appears more generic and robust.
- Our model relaxes strong assumptions. In [11], the au-

thors assumed that the probability of observing a feature vector $([lon_t, lat_t, SOG_t, COG_t]^T$ in their case) of a vessel at the time t , given its position and assigned route was independent. This assumption neglects the fact that AIS streams provide sequential data, feature vectors of a vessel's track are related to this vessel and interdependent. For instance, for such approaches the two branches of the "U" in Fig. 7b are normal.

- Methods in the literature do not deal with irregularly time-sampling problem. For example, model in [11] used sliding window to avoid incomplete tracks, and processed only the most recent points of the partially observed tracks. The vessel in Fig. 7b can outsmart this model by switching off its AIS transponder when performing the U-turn (which lasts about 30 minutes).
- In a complicated region like the Gulf of Mexico, all the methods based on DBSCAN [7], [11] cannot apply since

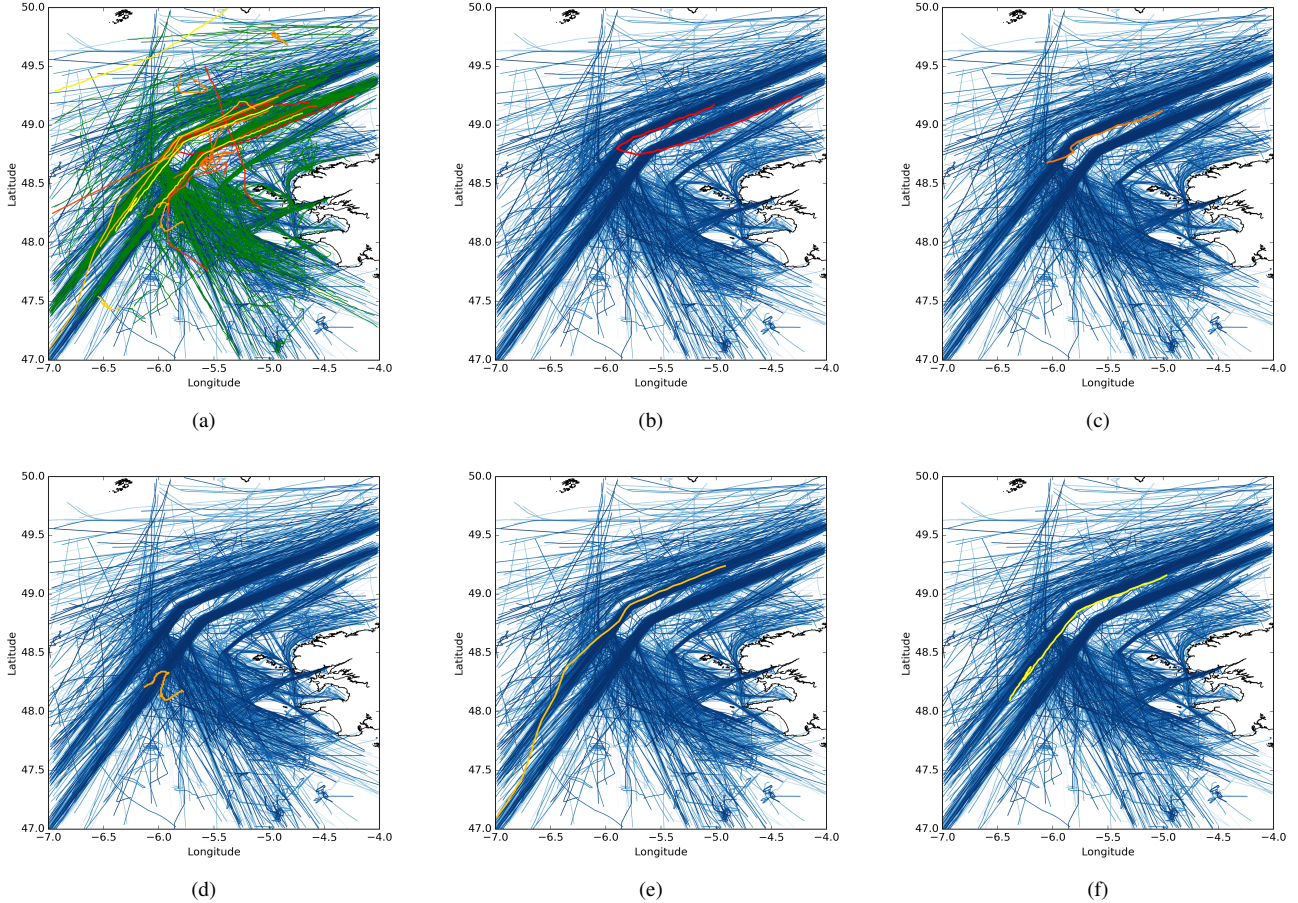


Fig. 7. Abnormal tracks detected by the proposed *a contrario* model (Brittany dataset). (a) All tracks detected in the test set; blue: tracks in the training set; green: normal tracks in the test set; other colors: abnormal tracks in the test set. (b) Abnormal U-turn. (c-d) Divergences from maritime route. (e) Abnormal route change. (f) Abnormal double-U-turn.

DBSCAN fails to extract effective waypoints. As shown in Fig. 5, in this area, vessels do not enter or exit the ROI at some specific zones, in consequence, DBSCAN can not detect entry and exist waypoints; beside that, a lot of vessels stop at sea for purposes (fishing for example), leads to false stationary waypoint detection by DBSCAN.

E. Vessel type identification

We tested the Vessel type identification submodel with a set of 1800 AIS tracks of 4 types of vessels: cargo, passenger, tanker and tug.

We compared the performance of our model with the one of other types of neural networks: CNNs and LSTMs (which are currently the state-of-the-art for time series classification). To simulate the missing data phenomenon in AIS streams, we deleted a 2-hours segment in each AIS track. Constant velocity model was used to fill the missing points for CNN model. We tested the LSTM networks with and without the “four-hot encoding” layer to show the benefit of this presentation. For each type of architecture we tried several configurations and report the best result.

The results are shown in Table. II. First, the poor performance of LSTMs without the “four-hot encoding” layer shows the relevance of this presentation for disentangling the explanatory information in continuous feature spaces of AIS messages’ attributes. Second, we can see that the proposed model achieved comparable performances with those of the state-of-the-art methods. It is because the embedding layer can provide a solid regular series of hidden regimes despite irregular time sampling in AIS streams. In addition to the slight improvement of the classification performance (from 87.43% to 87.72%), the proposed model also significantly reduces storage redundancies and computational requirements when doing each task separately, which is highly beneficial in an AIS big data context.

VI. INSIGHTS ON THE CONSIDERED APPROACH

In this section, we further discuss the key features of the considered approach with respect to state-of-the-art approaches. Overall, AIS vessel tracks (and trajectory data in general) may be characterized according to the following features:

TABLE II
CLASSIFICATION RESULTS.

Model	Precision	Recall	F1-score
LSTM	47.51%	64.11%	52.08%
LSTM_4-hot	88.04%	87.16%	87.43%
CNN	83.83%	84.06%	83.75%
<i>VRNN-CNN</i>	<i>88.00%</i>	<i>87.67%</i>	<i>87.72%</i>

- Time evolutions in terms of vessel position, speed and course;
- Geographical patterns (where is the vessel?);
- Geometric patterns (what is the shape of the track?);
- Speed and course distributions;
- Spatial-temporal patterns, called “*phase patterns*” (moves fast in specific zones and slowly in other zones for example).

We discuss below how our approach addresses the learning of these key features.

In time series, different features change at different temporal and spatial scales [25]. The proposed model learns these features from different points of view at different scales. At micro-scales, it learns the evolutions of the trajectories, *e.g.* with this historical information, in 10 minutes, vessel “V” seems to appear in zone “Z”, maintain its speed around “S” knots. These evolutions are modeled by the distribution $p(x_t|x_{1:t-1}, z_{1:t-1})$. At macro-scales, the model tends to learn the patterns of the entire AIS tracks.

Viewing the “four-hot” representation as an image-based representation of a track seems relevant to understand how our model can learn complex space-time patterns. More precisely, the one-hot vectors of the latitude/longitude coordinates of AIS messages indicate the rows/columns of the pixels in the image, respectively. Hence, if we cumulate these two one-hot vectors over a given time period, we build an image-based representation, which describes the geometric pattern of the vessel track (Fig. 8).

The proposed model is not translation-invariant and can learn spatial patterns and the geographical distributions of vessel tracks, *i.e.* a given type of tracks should appear in zone “A” and not in an other zone “B”. The phase patterns, on the other hand, reflect the correlations between temporal features and spatial ones. One typical example of phase pattern in trajectory data is the speed-position correlation, *e.g.* the average speed of vehicles on highway is higher than the one in urban area. Methods that use only the positions (longitude and latitude coordinates) to model trajectory and consider the speed as the first-order derivative of the positions can not capture this information. For example, the two tracks depicted in Fig. 9 are two examples from the processed dataset. They are similar in terms of spatial patterns, but different in terms of phase patterns, space-speed time series. Despite inter-individual variabilities, these two tracks exhibit in some regions low vessel speed and high vessel speed in other regions.




Fig. 8. Geometric patterns appear by summing up one-hot vectors of latitude and longitude coordinates.




Fig. 9. Illustration of phase patterns. We report the two examples of two AIS tracks of the processed dataset (red and blue). The solid lines are 3D curves (latitude, longitude and speed time series) reflect the phase pattern, whereas the dash curves (latitude and longitude time series only) reflect the associated 2D geometric patterns, which can not reveal the observed phase patterns.

These different aspects are similar to the wave-particle duality in physics, where the patterns correspond to the wave properties and the evolutions correspond to the particle properties.

VII. CONCLUSIONS AND PERSPECTIVES

In this paper, we proposed a novel deep-learning-based scheme for maritime traffic surveillance from AIS data streams. Stated within a probabilistic framework using Variational RNN, our approach overcomes strong limitations of state-of-the-art methods to jointly address multi-task issues, namely abnormal behavior detection, trajectory reconstruction and vessel type identification, on a regional scale, that is to say for datasets of spatially-heterogeneous datasets of tens or hundreds of millions of AIS data. More precisely, we tackled three main drawbacks of state-of-the-art approaches:

- First, we relax strong assumptions usually considered such as a finite number of behavioral categories (or hidden regimes) [12], [13].
- Second, by using VRNN, we can capture the maritime contextual information while avoiding problems that may be encountered if doing clustering.
- Third, the Embedding block in our model can deal with noise and irregularly time-sampling of AIS data streams. Besides, the Embedding block also results in an efficient compression of the behavioral information conveyed in data, which avoids making accessible the entire training

dataset for the operational use of the trained model. This appears critical for an operational big-data-compliant AIS system.

We also discussed the key aspects of the considered trajectory data representation, which is embedded in the considered VRNN framework.

Beyond benchmarking issues for large-scale datasets, including the evaluation of the ability of the proposed approach to scale up to global AIS data streams, the fusion with other sources of information available in the maritime domain could be a promising solution. Weather and ocean conditions, such as sea surface winds and currents, are two important factors that effect the behaviors of vessels. The exploitation of such variables should further constrain the considered VRNN framework and improve its representativity. The inference of behavioral models in low-density areas might require specific investigations in future studies, for instance some type of regularization.

VIII. ACKNOWLEDGEMENTS

This work was supported by public funds (Ministère de l'Éducation Nationale, de l'Enseignement Supérieur et de la Recherche, FEDER, Région Bretagne, Conseil Général du Finistère, Brest Métropole) and by Institut Mines Télécom, received in the framework of the VIGISAT program managed by "Groupement Bretagne Télédétection" (BreTel).

The authors acknowledge the support of DGA (Direction Générale de l'Armement) and ANR (French Agence Nationale de la Recherche) under reference ANR-16-ASTR-0026 (SESAME initiative), the labex Cominlabs, the Brittany Council and the GIS BRETEL (CPER/FEDER framework).

We also would like to thank MarineCadastre for the Gulf of Mexico AIS dataset and Collecte Localisation Satellites as well as Erwan Guegueniat for the Brittany AIS dataset.

REFERENCES

- [1] N. Perobelli, "MarineTraffic-A day in numbers," Jun. 2016. [Online]. Available: <https://www.marinetraffic.com/blog/a-day-in-numbers/>
- [2] R. A. Best and J. P. Norton, "A new model and efficient tracker for a target with curvilinear motion," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 3, pp. 1030–1037, Jul. 1997.
- [3] L. P. Perera, P. Oliveira, and C. G. Soares, "Maritime Traffic Monitoring Based on Vessel Detection, Tracking, State Estimation, and Trajectory Prediction," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 3, pp. 1188–1200, Sep. 2012.
- [4] R. Schubert, E. Richter, and G. Wanielik, "Comparison and evaluation of advanced motion models for vehicle tracking," in *2008 11th International Conference on Information Fusion*, Jun. 2008, pp. 1–6.
- [5] L. M. Millefiori, P. Braca, K. Bryan, and P. Willett, "Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 5, pp. 2313–2330, Oct. 2016.
- [6] G. Pallotta, S. Horn, P. Braca, and K. Bryan, "Context-enhanced vessel prediction based on Ornstein-Uhlenbeck processes using historical AIS traffic patterns: Real-world experimental results," in *17th International Conference on Information Fusion (FUSION)*, Jul. 2014, pp. 1–7.
- [7] F. Mazzarella, V. F. Arguedas, and M. Vespe, "Knowledge-based vessel position prediction using historical AIS data," in *2015 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, Oct. 2015, pp. 1–6.
- [8] S. Hexeberg, A. L. Flten, B. O. H. Eriksen, and E. F. Brekke, "AIS-based vessel trajectory prediction," in *2017 20th International Conference on Information Fusion (Fusion)*, Jul. 2017, pp. 1–8.
- [9] P. Coscia, P. Braca, L. M. Millefiori, F. A. N. Palmieri, and P. Willett, "Multiple Ornstein-Uhlenbeck Processes for Maritime Traffic Graph Representation," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2018.
- [10] J.-G. Lee, J. Han, and K.-Y. Whang, "Trajectory Clustering: A Partition-and-group Framework," in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '07. New York, NY, USA: ACM, 2007, pp. 593–604.
- [11] 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, Jun. 2013.
- [12] A. Holst, P. Ryman, and A. Linse, "Statistical Anomaly Detection for Maritime Surveillance and Monitoring," in *Maritime Knowledge Discovery and Anomaly Detection Workshop*, Joint Research Centre, ISPRA, Italy, Jul. 2016.
- [13] P. Gaspar, R. Lopez, M. Marzuki, R. Fablet, P. Gros, J.-M. Zigna, and G. Fabbrius, "Analysis of Vessel Trajectories for Maritime Surveillance and Fisheries Management," in *Maritime Knowledge Discovery and Anomaly Detection Workshop*, Joint Research Centre, ISPRA, Italy, Jul. 2016.
- [14] B. J. Rhodes, N. A. Bomberger, M. Seibert, and A. M. Waxman, "Maritime situation monitoring and awareness using learning mechanisms," in *MILCOM 2005 - 2005 IEEE Military Communications Conference*, Oct. 2005, pp. 646–652 Vol. 1.
- [15] R. Laxhammar, "Anomaly detection for sea surveillance," in *2008 11th International Conference on Information Fusion*, Jun. 2008, pp. 1–8.
- [16] B. Ristic, B. L. Scala, M. Morelande, and N. Gordon, "Statistical analysis of motion patterns in AIS Data: Anomaly detection and motion prediction," in *2008 11th International Conference on Information Fusion*, Jun. 2008, pp. 1–7.
- [17] K. Kowalska and L. Peel, "Maritime anomaly detection using Gaussian Process active learning," in *2012 15th International Conference on Information Fusion*, Jul. 2012, pp. 1164–1171.
- [18] J. Will, L. Peel, and C. Claxton, "Fast Maritime Anomaly Detection using KD Tree Gaussian Processes," in *2nd IMA Conference on Maths in Defence*, 2011.
- [19] F. Johansson and G. Falkman, "Detection of vessel anomalies - a Bayesian network approach," in *Sensor Networks and Information 2007 3rd International Conference on Intelligent Sensors*, Dec. 2007, pp. 395–400.
- [20] S. Mascaro, A. E. Nicholso, and K. B. Korb, "Anomaly detection in vessel tracks using Bayesian networks," *International Journal of Approximate Reasoning*, vol. 55, no. 1, Part 1, pp. 84–98, Jan. 2014.
- [21] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015. [Online]. Available: <https://www.nature.com/articles/nature14539>
- [22] J. Chung, K. Kastner, L. Dinh, K. Goel, A. Courville, and Y. Bengio, "A Recurrent Latent Variable Model for Sequential Data," in *Advances in neural information processing systems*, Jun. 2015, pp. 2980–2988.
- [23] M. Fraccaro, S. r. K. S nderby, U. Paquet, and O. Winther, "Sequential Neural Models with Stochastic Layers," in *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2016, pp. 2199–2207.
- [24] C. J. Maddison, D. Lawson, G. Tucker, N. Heess, M. Norouzi, A. Mnih, A. Doucet, and Y. W. Teh, "Filtering Variational Objectives," in *Advances in Neural Information Processing Systems*, May 2017, pp. 6576–6586.
- [25] Y. Bengio, A. Courville, and P. Vincent, "Representation Learning: A Review and New Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
- [26] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *arXiv:1312.6114 [cs, stat]*, Dec. 2013, arXiv: 1312.6114. [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [27] X. Jiang, E. N. de Souza, A. Pesaranghader, B. Hu, D. L. Silver, and S. Matwin, "TrajectoryNet: An Embedded GPS Trajectory Representation for Point-based Classification Using Recurrent Neural Networks," *arXiv:1705.02636 [cs]*, May 2017, arXiv: 1705.02636. [Online]. Available: <http://arxiv.org/abs/1705.02636>
- [28] M. Ammar and S. Le Hegarat-Masclé, "An A-Contrario Approach for Object Detection in Video Sequence," *International Journal of Pure and Applied Mathematics*, vol. 89, no. 2, Dec. 2013.
- [29] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.