Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)
Sister Conferences Best Papers Track

Logic Rules Meet Deep Learning: A Novel Approach for Ship Type Classification (Extended Abstract)*

Manolis Pitsikalis[1], **Thanh-Toan Do**[2], **Alexei Lisitsa**[1] and **Shan Luo**[1,3]
[1]Department of Computer Science, University of Liverpool
[2]Department of Data Science and AI, Monash University
[3]Department of Engineering, King's College London
{e.pitsikalis,a.lisitsa,shan.luo}@liverpool.ac.uk
toan.do@monash.edu

Abstract

The shipping industry is an important component of the global trade and economy. In order to ensure law compliance and safety, it needs to be monitored. In this paper, we present a novel ship type classification model that combines vessel transmitted data from the Automatic Identification System, with vessel imagery. The main components of our approach are the Faster R-CNN Deep Neural Network and a Neuro-Fuzzy system with IF-THEN rules. We evaluate our model using real world data and showcase the advantages of this combination while also compare it with other methods. Results show that our model can increase prediction scores by up to 15.4% when compared with the next best model we considered, while also maintaining a level of explainability as opposed to common black box approaches.

1 Introduction

Nowadays, the combination of deep learning with logic has been attracting a lot of attention for numerous reasons. On the one hand, deep learning has been used extensively in different applications, such as object detection [Ren *et al.*, 2015; Redmon *et al.*, 2016] and language analytics tasks [Mikolov *et al.*, 2013; Collobert and Weston, 2008] with a lot of success. On the other hand, logic based approaches have been used widely in tasks where explainability is required, such is the case in certain medical tasks [London, 2019], or where expert knowledge is available and needs to be encoded into a model [Grosan and Abraham, 2011]. However, logic based approaches typically provide crisp predictions over symbolic data, and it is often the case that human experts are required to express the knowledge into some form of logic. At the same time, although, deep learning approaches can handle unstructured data such as text and images, their black box nature makes them inadequate for tasks where explainability is a key requirement. For these reasons, there is a need for a model that combines the advantages of deep learning with those of logic based approaches. A notable example from this direction, is the approach of [Hu *et al.*, 2016], where knowledge from logic rules is harnessed via a teacher-student model setting. While the approach of [Hu *et al.*, 2016] manages to improve accuracy in certain Natural Language Analytics tasks, it does not allow any learning in the aspect of the logic rules as they remain constant during training. Another work that overcomes the crisp values of logic rules is the approach of [Tsipouras *et al.*, 2008]. Their approach, similar to ours, involves the fuzzification of logic rules extracted from a C4.5 decision tree for the diagnosis of coronary artery disease. In this paper, we propose a model that combines deep learning with fuzzified learned logic rules, applied in the maritime domain, for the task of ship type classification.

Shipping is a pillar of the global trade and economy, but there are many cases where ships are found to be involved in illegal activities. Fortunately, there is an abundance of maritime data sources that can provide valuable information for maritime surveillance and in our case ship type classification. When it comes to ship types, maritime vessels are divided based on their characteristics and purpose, e.g., fishing vessels, cargo vessels etc. Consequently, different regulations apply to each ship type. To promote security and abidance to regulations the process of classifying and validating a vessel's type needs to be automated using the data available.

In this paper, we perform ship type classification using two data sources. We use static vessel transmitted data from the Automatic Identification System (AIS) that allows the transmission of dynamic spatio-temporal data and static identity data from vessels. Moreover, since cameras are widely used in maritime settings (e.g., CCTV in ports), we use additionally a dataset of images pre-linked with the vessels existing in the AIS records. The main components of our approach are the Faster R-CNN deep neural network and a Neuro-Fuzzy model leveraging convolutional deep features of the former along with AIS information. We evaluate the presented model using real world data, against other algorithms available. A larger version of this paper is available in [Pitsikalis *et al.*, 2021].

2 Methodology

The main components of our approach are first a Fuzzy model created by extracting and fuzzifying rules from a set of Classi-

*This is an extended abstract of a paper that appeared/won the "Harold Boley award for most promising paper" at the RuleML+RR 2021 conference.

fication and Regression Trees (CART) [Breiman *et al.*, 1984], and second the Faster R-CNN deep neural network [Ren *et al.*, 2015], a deep neural network used for object detection. In this Section, we present in detail the methodology for rule extraction and fuzzification, deep feature extraction from the Faster R-CNN deep neural network and finally the Neuro-Fuzzy Combination capable of handling two data information sources—the first source contains deep features from images while the second one contains numerical structured data describing attributes of vessels—with the aim achieving higher prediction scores and increased explainability.

2.1 Rule Extraction and Fuzzification

The first stage of our methodology involves the extraction of logic rules in Disjunctional Normal Form (DNF). For each class label y, we train a Decision Tree model using the CART algorithm for binary classification where the positive class is y and the remaining are negative. Then, for each class y we parse the corresponding $tree$ and recursively create a rule by adding conditions expressing the path from the root to leaf nodes where the label is y. Therefore, a condition C_i included in the body of the rule concerning label y is expressed as:

$$C_i = (x_1 \; op \; v_1) \wedge \cdots \wedge (x_k \; op \; v_k) \tag{1}$$

where v_i are the values obtained during the splitting process of the tree training; x_i are the values of the attributes on which the comparisons are applied and op is either '$>$' or '$\leq$'. A rule R_i for a specific class y_i has the following form:

$$IF \; C_1 \vee \cdots \vee C_n \; THEN \; y_i \tag{2}$$

For the fuzzification of the rules, we apply the sigmoidal membership function (3) in the comparisons $c_{ij} \in C_i$ so that each fuzzified comparison c'_{ij} yields a value in $(0, 1)$:

$$c'_{>,\leq}(x; s, v) = \frac{1}{1 + e^{-\kappa s(x-v)}}, \; \kappa \in \{1, -1\} \tag{3}$$

where s is the slope of the sigmoid curve, v is the 'center' of the curve and κ is 1 or -1 if the comparison is $x > v$ or $x \leq v$ respectively (see Figure 1).

Moreover, as seen in Equation (4), for each class rule R_i we multiply each of the conditions C_i with a weight w_i

$$IF \; w_1 C_1 \vee \cdots \vee w_n C_n \; THEN \; y_i \tag{4}$$

where $w_i \in [0, 1]$ and $\sum_{i=0}^{n} w_i = 1$.

Finally, we replace the logical connectives of conjunction and disjunction with their Weighted Exponential Mean approximation [Dujmović and Larsen, 2007] as depicted in Table 1, $\min'$ and $\max'$ respectively. For the experimental results presented in Section 3 we have set the level of andness and orness to 'medium high' $r = -5.4$ for conjunctions and $r = 5.4$ for disjunction by taking into consideration the results of the ablation study presented in [Pitsikalis *et al.*, 2021]. Consequently, the truth value of a rule R_i is produced by computing the $\max'$ of all the conditions C_i as follows:

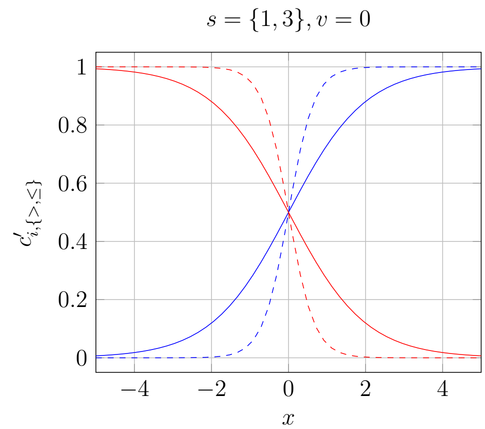

Figure 1: Plots of $c'_{i,>}$ (blue) and $c'_{i,\leq}$ (red) with s set to 1 and 3 (continuous and dashed lines respectively) for $x \in [-5, 5]$.

Notation	WEM
$\bigwedge_{i=1}^{n} c_i$	$\frac{1}{r} \ln \left(\frac{1}{n} \sum_{i=0}^{n} e^{r c_i} \right), r \in \{-14.0, -5.4, -2.14\}$
$\bigvee_{i=1}^{n} w_i C_i$	$\frac{1}{r} \ln \left(\sum_{i=0}^{n} w_i e^{r C_i} \right), r \in \{2.14, 5.4, 14.0\}$

Table 1: Approximation of conjunction and disjunction using Weighted Exponential Means. Note that in the case of conjunction, equal weights are used. An analysis of the effects of r for WEM can be found in [Dujmović and Larsen, 2007].

$$R_i(x; W, S) = \max'_{i \in [1,n]} \left\{ w_i \min'_{j \in [1, |C_i|]} c'_{ij}(x_{ij}; s_{ij}, v_{ij}) \right\} \tag{5}$$

where W is a vector containing the weights of the conditions, S is a vector containing the slope parameters of the sigmoid membership functions included in the fuzzified comparisons of each C_i while $|C_i|$ is the number of fuzzified comparisons included in C_i. Finally, to make a class prediction, we produce the R_i values for each class y_i and produce the simplex vector by applying $L1$ normalisation on the vector $F_c = \{R_1, \ldots R_m\}$ where m is the number of classes.

2.2 Neuro-Fuzzy Combination

In the previous Section, we mentioned that each class rule R_i accepts as parameters a set of weights W for the weighted approximation of disjunction and a set of slopes S for the sigmoid membership functions included in the conditions of the rule. Here, we present how we combine the Fuzzy Model described in Section 2.1 with a Neural Network into a single model. The architecture of our combined model is illustrated in Figure 2a.

We perform ship type classification by combining the information included in the images and the AIS transmitted characteristics of vessels. For this reason, the architecture of our model has two inputs. The first one accepts deep convolutional features corresponding to images, while the second input includes the values of the AIS fields. In detail, for each

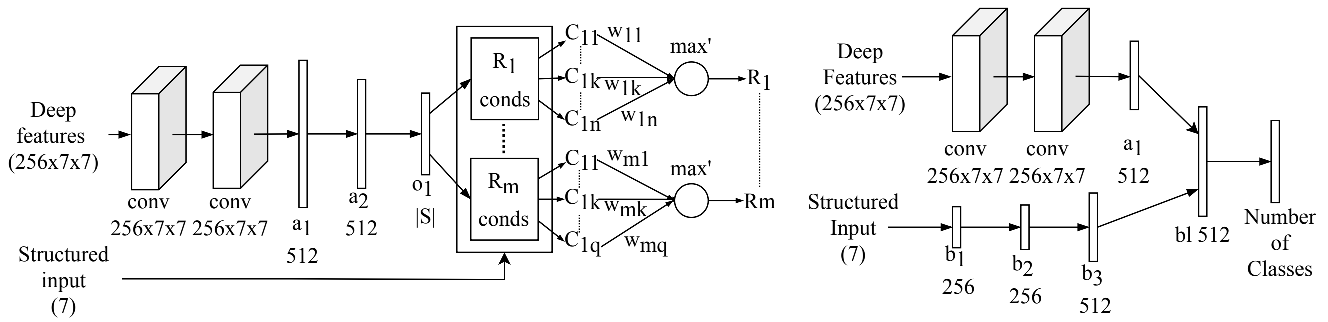

(a) Our Neuro-Fuzzy network

(b) Baseline model

Figure 2: (a) Architecture of the Neuro-Fuzzy network. The upper branch receives deep features extracted from the RoI pooling layer of a Faster R-CNN model when given a vessel image, while the lower branch receives AIS data (Structured data) that is then given to the different fuzzified rules $R_{i\ conds}$. (b) Architecture of the baseline model. Here, in contrast with the model of (a), AIS data is given as input to the b_1 layer of the lower branch, while the combination of the different data sources is achieved using the bilinear layer bl.

vessel image we extract a deep convolutional feature from a pre-trained Faster R-CNN network in prediction mode. We keep the $256 \times 7 \times 7$ feature vector from the output of the RoI pooling layer, corresponding to the bounding box that yields the highest confidence score after the non-maximum suppression stage. In the upper branch of Figure 2a we use two convolutional layers, where the output of the second layer is flattened and given as input to a fully connected layer a_1 with 512 neurons followed by the fully connected layer a_2, with batch normalization, dropout and ReLu activation, and finally the o_1 layer which has Leaky ReLu as an activation function and yields an output equal to the size of S, i.e, the vector containing the slope parameters of the rules.

Then, using the output of the o_1 layer, along with the rules input (AIS fields) we can now compute the values C_{ij} included in each $R_i, i \in [1, m]$ (see the $R_{i\ conds}$ blocks in Figure 2a). Next, for each $R_{i\ conds}$ we create the vector $\{e^{rC_{i1}}, \ldots, e^{rC_{in}}\}$ and feed it into a log activated layer, with bias set to 0 and normalised weights, that computes the approximation of weighted disjunction as follows:

$$R_i = \frac{1}{r} \ln \left(\sum_{j=0}^{n} w_{ij} e^{rC_{ij}} \right) \tag{6}$$

where w_{ij} are the weights of the input and r is the orness level. Finally, all R_i are fed through a *softmax* layer that outputs a probabilistic vector F. We train the complete neuro-fuzzy model using the cross entropy loss L over F and the one-hot ground truth label $y = \{y_c\}_1^m$:

$$L = -\sum_{c=1}^{m} y_c \log(f_c),\ F = \{f_1, \ldots f_m\} \tag{7}$$

3 Evaluation

In this section, we present the characteristics of the datasets we use for our experimental evaluation, the experimental settings for the training of our models and finally the prediction scores of the evaluated models.

AIS Field	Description
MMSI	Maritime Mobile Service Identity
to_{bow, stern, starboard, port}	Distance from the AIS transceiver to the vessel's {bow, stern, right side, left side}
draught	Vertical distance between the waterline and the bottom of the vessel's hull
width	to bow + to stern (vessel's width)
length	to starboard + to port (vessel's length)

Table 2: Description of the retained AIS fields.

3.1 Dataset

We construct our dataset using the AIS records from the maritime dataset presented in [Ray *et al.*, 2019]. In the context of the presented experiments, we use only the fields presented in Table 2, of the static AIS messages. Moreover, using the 'MMSI' value we collect up to 5 images for each vessel from the IHS Markit World Register of Ships (v12) and the ShipScape photographic library. The collected images contain one vessel per image and have been annotated using the ship type field of the AIS messages and the manual selection of the bounding box of each ship. Using the retained AIS fields and the collected images we create an Image classification centred dataset IC containing a deep feature for each image and the AIS fields corresponding to the vessel in the image. However, there is also another way of looking at the classification problem. While in the previous case the problem is image centralised, in the current case we focus on the vessels, therefore we create a vessel centred dataset VC by grouping and averaging the deep features per vessel MMSI. Therefore, VC contains an AIS record for each vessel, and a deep feature created by averaging the deep features corresponding to images of that vessel. The number of vessels and images per vessel type are presented in the left part of Table 3. Note that vessels that don't have images available, and ship types that have less than 20 different vessels have been omitted from the experiments and are not included in Table 3.

(a) Cargo. (b) Tanker. (c) Tug. (d) Passenger. (e) Other.

Figure 3: Example detections. The bounding boxes are produced by the Faster R-CNN network while the detected labels are produced from the Neuro-Fuzzy model of this paper.

Dataset Characteristics			Image Centred (mAP)			Vessel Centred (Macro F1-Score)							
Ship type	Vessels	Images	OM	B	FRCNN*	OM	B	DT*	kNN*	NB*	LR*	LDA*	SVM*
Cargo	2412	11185	91	92	**94**	**97**	**97**	88	91	74	86	84	82
Tanker	864	3950	87	84	**93**	94	**95**	72	71	46	46	38	10
Other	53	229	**9**	**9**	**9**	**37**	23	27	30	0	22	0	0
Passenger	42	199	94	**97**	84	88	**94**	37	22	46	0	28	0
Tug	32	139	**67**	59	60	88	18	80	**89**	82	40	22	88
All	3403	15702	**69.6**	68.2	68.0	**80.8**	65.4	60.8	60.6	49.6	38.8	34.4	36.0

Table 3: Dataset characteristics (left of the double vertical line). mAP and Macro F1-Scores of the evaluated models (right of the double vertical line). OM, FRCNN, B, DT, kNN, NB, LR, LDA and SVM stand for 'Our Model', 'Faster R-CNN', 'Baseline', 'Decision Tree', 'k Nearest Neighbours', 'Logistic Regression', 'Linear Discriminant Analysis' and 'Support Vector Machines'. Bold values indicate the highest score per ship and dataset type. The confidence threshold of retaining a bounding box, during the prediction phase of Faster R-CNN, has been set has been set to 0.7. Models with an '*' used only one source of information i.e., either Images or AIS records.

3.2 Baseline Model

In addition to the Neuro-Fuzzy model presented in this paper, we create the baseline model of Figure 2b which retains the convolutional branch of the neuro-fuzzy model up to layer a_1 and adds a second branch that accepts as input the AIS fields. The additional branch has one input layer with 7 input neurons and 256 output neurons, followed by two fully connected layers (b_2, b_3) with batch normalization, dropout and ReLu activation. The output of layers a_1 and b_3 is then given as input to the bilinear layer bl, which has batch normalization, dropout and ReLu activation. Finally, the output of layer bl is fed into a fully connected output layer with $softmax$ activation that yields the class prediction.

3.3 Experimental Setup

We extract the IF-THEN rules of the Neuro-Fuzzy model using 75% of the AIS records (minus the 'MMSI' field) and train a Faster R-CNN model, with 'ResNet-50' [He *et al.*, 2016] network as backbone using the corresponding vessel images. Then, we extract the deep feature corresponding to each image and create the datasets IC and VC. Finally, we train both our Neuro-Fuzzy model and the baseline model on the created datasets IC and VC and evaluate them separately.

3.4 Experimental Results

We evaluate our model on both Image and Vessel centred datasets. In the first case we use the mean Average Precision Metric (mAP) presented in [Padilla *et al.*, 2020], with interpolation over all recall levels while in the second case we use the macro F1-Score. Some example detections of the Neuro-Fuzzy model are illustated in Figure 3. The right part of Table 3 shows that the combination of vessel transmitted AIS information along with Imagery using the Neuro-Fuzzy model of this paper yields better results than using each data source separately and using both sources in the baseline model in both Image and Vessel centred datasets. However, although data fusion proves to improve prediction scores, we attribute the low prediction scores to the class imbalance of the dataset, since the lowest mAP and F1 scores where produced by the 'Other' ship type which expresses a diverse spectrum of vessels but has very few examples in the present case. Moreover, in the image centred dataset, although the score difference is not significant, our model compared to the other two offers to some degree explainability since a classification decision can be tracked through the rules included in the Neuro-Fuzzy system. An ablation study and an example of an extracted rule along with its fuzzified version can be found in [Pitsikalis *et al.*, 2021].

4 Conclusions and Future Work

We presented a methodology that can be used to combine effectively AIS data with vessel Imagery for ship type classification. We believe that the logic rules extracted by the decision trees add information over the dependencies between the AIS fields, and thus providing additional information in the combined Neuro-Fuzzy model. Although our methodology has been applied in the maritime domain, we believe that it can be also applied in other domains where multiple sources of information are available.

Acknowledgements

This work has been funded by the EPSRC Centre for Doctoral Training in Distributed Algorithms at the University of Liverpool, and Denbridge Marine Limited[1], United Kingdom.

References

[Breiman *et al.*, 1984] Leo Breiman, Jerome Friedman, R.A. Olshen, and Charles J. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA, 1984.

[Collobert and Weston, 2008] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, page 160–167, New York, NY, USA, 2008. Association for Computing Machinery.

[Dujmović and Larsen, 2007] Jozo J. Dujmović and Henrik Legind Larsen. Generalized conjunction/disjunction. *International Journal of Approximate Reasoning*, 46(3):423–446, 2007. Special Section: Aggregation Operators.

[Grosan and Abraham, 2011] Crina Grosan and Ajith Abraham. *Rule-Based Expert Systems*, pages 149–185. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

[Hu *et al.*, 2016] Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. Harnessing deep neural networks with logic rules. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2410–2420, Berlin, Germany, August 2016. Association for Computational Linguistics.

[London, 2019] Alex John London. Artificial intelligence and black-box medical decisions: Accuracy versus explainability. *Hastings Center Report*, 49(1):15–21, 2019.

[Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, page 3111–3119, Red Hook, NY, USA, 2013. Curran Associates Inc.

[Padilla *et al.*, 2020] Rafael Padilla, Sergio L. Netto, and Eduardo A. B. da Silva. A survey on performance metrics for object-detection algorithms. In *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*, pages 237–242, 2020.

[Pitsikalis *et al.*, 2021] Manolis Pitsikalis, Thanh-Toan Do, Alexei Lisitsa, and Shan Luo. Logic rules meet deep learning: A novel approach for ship type classification. In Sotiris Moschoyiannis, Rafael Peñaloza, Jan Vanthienen, Ahmet Soylu, and Dumitru Roman, editors, *Rules and Reasoning*, pages 203–217, Cham, 2021. Springer International Publishing.

[1]https://www.denbridgemarine.com

[Ray *et al.*, 2019] Cyril Ray, Richard Dréo, Elena Camossi, Anne-Laure Jousselme, and Clément Iphar. Heterogeneous integrated dataset for maritime intelligence, surveillance, and reconnaissance. *Data in Brief*, 25:104141, 2019.

[Redmon *et al.*, 2016] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, Los Alamitos, CA, USA, jun 2016. IEEE Computer Society.

[Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 91–99, Cambridge, MA, USA, 2015. MIT Press.

[Tsipouras *et al.*, 2008] Markos G. Tsipouras, Themis P. Exarchos, Dimitrios I. Fotiadis, Anna P. Kotsia, Konstantinos V. Vakalis, Katerina K. Naka, and Lampros K. Michalis. Automated diagnosis of coronary artery disease based on data mining and fuzzy modeling. *IEEE Transactions on Information Technology in Biomedicine*, 12(4):447–458, 2008.