

Robust Uncertainty Estimation for Classification of Maritime Objects

Jonathan Beckett, Frederik Schöller, Evangelos Boukas, and Lazaros Nalpantidis

Abstract—We explore the use of uncertainty estimation in the maritime domain, showing the efficacy on toy datasets (CIFAR10) and proving it on an in-house dataset, SHIPS. We present a method joining the intra-class uncertainty achieved using Monte Carlo Dropout, with recent discoveries in the field of outlier detection, to gain more holistic uncertainty measures. We explore the relationship between the introduced uncertainty measures and examine how well they work on CIFAR10 and in a real-life setting. Our work improves the FPR95 by 8% compared to the current highest-performing work when the models are trained without out-of-distribution data. We increase the performance by 77% compared to a vanilla implementation of the Wide ResNet. We release the SHIPS dataset and show the effectiveness of our method by improving the FPR95 by 44.2% with respect to the baseline. Our approach is model agnostic, easy to implement, and often does not require model retraining.

I. INTRODUCTION

The autonomous operation of robots, including autonomous ships, heavily relies on the perception of the world around them. Camera-based perception has seen a lot of growth in the past years. This is especially present in many well-known datasets, such as CIFAR10 and ImageNet, where models achieve almost perfect performance. However, this level of performance is often only achieved after careful reparameterization. The resulting models produce very accurate results on the specific datasets, but their performance degrades dramatically when applied to real-world robot operations, such as noisy and previously-unseen targets. Furthermore, the output of these models does not provide a usable measure of the certainty of their predictions, which is often required for robust autonomous operation. Even though the softmax function allows mapping of the logits into a probability distribution, this mapping is commonly not well calibrated, leading to over-confident predictions. While the softmax function also provides an intra-class measure of uncertainty, we will, in this work, present a method to classify samples and produce a more holistic uncertainty measure. We showcase this on a simple, commonly used dataset (CIFAR10) and on our own curated dataset (SHIPS) while also showing the performance against six outlier datasets.

All authors are with the Department of Electrical and Photonics Engineering, DTU - Technical University of Denmark `jbibe`, `fets`, `evanb`, `lanalpa@dtu.dk`

This research is sponsored by the Danish Innovation Fund, The Danish Maritime Fund, Orients Fund, and the Lauritzen Foundation through the Autonomy part of the ShippingLab project, Grant number 8090-00063B.

Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. ©2023 All rights reserved.



Fig. 1: The GreenHopper aims to autonomously traverse "Limfjorden" in Aalborg, Denmark. Photo: Facebook Port of Aalborg, GreenHopper inauguration (Nov. 2022).

Our primary focus is on autonomous operation at sea and, more precisely, on the predictive uncertainty for the classification of common maritime vessels and objects for the GreenHopper vessel, see Figure 1. In our previous work [1], [2], [3], we proposed an object detection network tasked with robust detection of two coarse classes, *buoys* and *ships*; given an image, a detection consisted of an object bounding box and class confidence. This work extends our efforts of creating a more reliable and robust object detection system [4], [5], by focusing on producing higher quality classification outputs, that is a more precise label (e.g. from boat to sailboat or motorboat) and providing a usable uncertainty metric for said classification.

A. Contributions

The main contributions of this paper are: (i) we propose a joint method to produce intra-class (aleatoric) and out-of-distribution (epistemic) uncertainty of our predictions. (ii) We show that our method, to the best of our knowledge, performs better than any other current work at outlier detection when only trained on ID data. (iii) We explore the relationship between the explored uncertainty measures and how they can be combined for better out-of-distribution detection. Finally, (iv) our work produces well-calibrated networks with usable uncertainty estimates.

II. RELATED WORK

Probabilistic inference has heavily inspired the work of network uncertainty estimates. Early work in this field, such as Bishop et al. [6], proposed using Mixture Density Networks (MDN) to estimate predictive distributions. The MDN models approximate the conditional distribution over a scalar response as a mixture of Gaussians. The parameters

of a Gaussian mixture describing the predictive distribution are estimated by training a model to output parameters maximizing the overall log-likelihood. The work of [7] proposed calibrating the output of neural networks by scaling the logits by a constant factor before the softmax function. It showed how a calibrated network could give a better probabilistic estimate of the likelihood of a prediction. Charpentier et al. [8] estimated the latent distribution of classes to detect out-of-distribution (OOD) examples. Training a model to output the parameters of a Dirichlet distribution, it was possible to estimate predictive uncertainty with a single forward pass and classification of OOD examples. The uncertainty estimate was used to improve segmentation results on brain scans. Hendrycks et al. [9] proposed to use the maximum softmax probability as a metric for outlier detection. The same group extended that work in [10] where they instead proposed using a subset of OOD datasets for the basis of OOD detection. Several methods for OOD detection have since been explored, such as energy score [11], where the authors propose to use the energy of the output vector for OOD detection. Du et al. [12] extend this by introducing a scheme that uses a multivariate distribution of a latent layer to create “virtual outliers”, which are then used for training. The work of Lee et al. introduced the Mahalanobis distance [13] as a measure for OOD. A method for detecting OOD examples in neural networks, dubbed ODIN, was introduced by Liang et al. in [14] which applies small perturbations to the input calibrated from an OOD sample. Hsu et al. expanded ODIN by introducing the generalized-ODIN [15], which proposes to decompose confidence scoring, removing the need to calibrate on OOD-Data. Further work, such as the generation and collection of outlier data to be used as regularization data for OOD detection, has been explored in several works, including [16], [10], [17]. In contrast, Grcic et al. proposed to train a generative model to synthesize outliers in the pixel space [18].

Producing usable uncertainty estimates for object classification is difficult; neural networks can produce a probability estimate by adding the softmax function. However, Guo et al. [19] highlights the importance of model calibration, where the miss-match in class confidence and the true positive rate is often skewed towards overconfidence. Kuleshov et al. [20] explores the use of model calibration to produce usable softmax probability estimates from networks. Lakshminarayanan et al. argue in [21] that using an ensemble of models can produce well-calibrated uncertainty estimates. The ensembles are generated by training multiple instances of a model on a random permutation of a given data set. The ensemble is treated as a uniformly-weighted mixture model, and the predictions are combined using the average of the outputs. Bayesian Neural Networks (BNN) [22], on the other hand, define all model parameters as a Gaussian distribution, with μ and Σ as the mean and covariance. When training a BNN, the model parameters are updated using Bayes theorem. In practical terms, this is done by minimizing the Kullback–Leibler divergence [23], where a predictive uncertainty is found by sampling the model weights.

III. BACKGROUND

The following section will outline the theory and methods used for our proposed work.

A. Epistemic and Aleatoric Uncertainty

In our work, we use the concepts of epistemic and aleatoric uncertainty [24] as two orthogonal uncertainty estimation metrics, inspired by Wang et al. [25].

Epistemic uncertainty refers to the uncertainty arising from a lack of knowledge. In machine learning, this occurs when our parameters provide an inadequate fit, often due to a lack of data, causing our posterior over parameters to be broad. Fig. 2 displays an example classification problem, where the middle plot captures epistemic uncertainty.

Aleatoric uncertainty, on the other hand, captures the stochasticity or variability in the data. Given a large dataset of high variance labels, the best possible prediction may be a high entropy one resulting in poor intra-class uncertainty. The example classification problem depicted in Fig. 2 shows the uncertainty between classes.

A more practical measure of uncertainty is the combination of epistemic and aleatoric uncertainty, as seen in the final plot of Fig. 2. In this case, the regions that the network has not encountered are identified as uncertain regions.

B. Monte Carlo Dropout

Dropout [26] was introduced to prevent overfitting in neural networks by disabling a percentage of randomly selected neurons during training. Each neuron has some probability of being disabled, called the dropout rate.

Monte Carlo Dropout (MC Dropout), proposed by [27], allows for the use of dropout as a Bayesian approximation of the Gaussian process. The network weights are alternately dropped and simulate Monte Carlo (MC) samples from the space of all available models. Each dropout configuration Θ_t corresponds to a different sample from the approximate parametric posterior distribution $Q(\Theta|D)$. Thus, sampling from the approximate posterior distribution $P(T|D)$ allows for MC sampling of the model likelihood.

C. Virtual Outlier Detection

Outlier estimation estimates whether a sample is in-distribution (ID) or out-of-distribution (OOD). Our approach for estimating ID and OOD samples is based on the virtual estimation synthesis (VOS) method presented in [12]. During training, a class conditional multivariate Gaussian for classes $c \in 1, \dots, C$ is created $P_\theta(h(\mathbf{x}) | y = c) = \mathcal{N}(\mu_c, \Sigma)$, where μ_c is the class mean and Σ is the covariance. By sampling the ϵ – likelihood on the class-conditional distribution, we sample outliers \mathcal{V}_k . These outliers are the basis for the work done in [12]. The outliers and the correct samples are minimized as a Binary classification problem. The goal is to reduce the energy for OOD samples and increase it for ID samples. The energy term is calculated as follows:

$$E(\mathbf{x}; \theta) = -\log \sum_{k=1}^K \exp^{f_k(\mathbf{x}; \theta)} \quad (1)$$

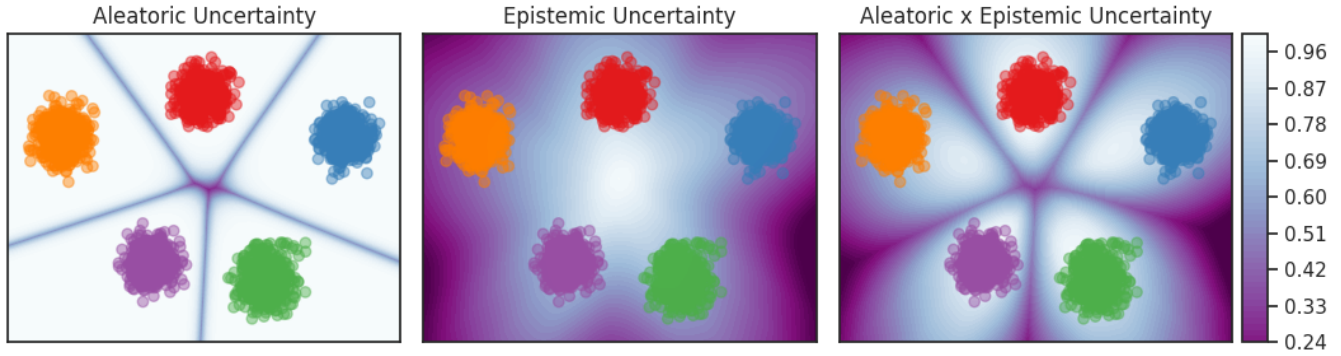


Fig. 2: The three plots present the probability map of the (x, y) feature space for 5 clusters. The first plot shows the intra-class (aleatoric) uncertainty, as seen in the center and between the clusters. The second plot provides the OOD (epistemic) uncertainty. The final plot shows the combination of the two providing a better uncertainty measure.

We scale the Energy term during training: $ES(\mathbf{x}; \theta) = \log\left(\frac{\mu_{c, ID}}{E(\mathbf{x}; \theta)}\right)$ to directly minimize the Energy with binary cross-entropy. To better control the VOS loss during training we compute the running mean for each class $\mu_{c, ID}$, given the respective energy $E(\mathbf{x}; \theta)$, resulting in better regulation of class energy inconsistencies. We use the binary cross-entropy loss (BCE) on the two meta-classes, positive samples (ID) and the negative samples (OOD), drawn from the created class-conditional multivariate distributions \hat{x} .

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathcal{L}_{\text{cls}} + \beta \cdot \mathcal{L}_{\text{uncert}}] \quad (2)$$

where \mathcal{L}_{cls} is the Cross-Entropy Loss and $\mathcal{L}_{\text{uncert}}$ is the BCE loss, which is scaled by a factor β .

D. Logit Normalization

Logit Normalization (LN) is introduced by Wei et.al [28] as a simple but effective method of reducing the number of outliers detected. This occurs by normalizing the predicted logits before computing the cross-validation loss. The loss is computed for a model f given as input a target pair (x, y) with weights θ as:

$$\mathcal{L}_{\text{ln}}(f(\mathbf{x}; \theta), y) = -\log \frac{e^{f_y / (\tau \|\mathbf{f}\|)}}{\sum_{i=1}^k e^{f_i / (\tau \|\mathbf{f}\|)}} \quad (3)$$

where τ is a temperature parameter that modulates the magnitude of the logits. The authors presented an impressive performance increase on an OOD dataset test suite commonly used to compare outlier detection works. The results are presented in Table III.

IV. METHODOLOGY

Our goal is to provide a well-calibrated model that produces epistemic and aleatoric uncertainty estimates. We propose to sample the VOS models with MC-Dropout; this would allow for the aleatoric uncertainty to be described by MC Dropout while the Epistemic uncertainty (the OOD samples) will be scored by the VOS Energy Score. Furthermore, with MC-Dropout, we should be able to handle extreme outliers better, as a set of Energy scores with a significant variance implies that the model is uncertain whether it is

OOD or ID. We propose using the MI of the samples from MC-Dropout as a supportive OOD detection measure, as the MI of the OOD datasets appears to be lower.

To further manage the Energy Score, we introduce the Logit Normalization loss see Eq. 3. By removing the incentive to increase the logits when computing the classification loss but allowing it when computing the loss on the Scaled energy, we aim to produce a network that better ensures that the magnitude of the energy is correlated to the state of a sample, whether it is ID or OOD. Furthermore, the introduction of Logit Normalization ensures calibrated networks.

A. Uncertainty measures

Following established practices in the literature, we use the following metrics to measure the performance of detections:

1) *Aleatoric uncertainty similarity measures:* Following recent work, we here present tools for quantifying the aleatoric uncertainty. The mutual information (MI) explored by Smith et al. [24] in a machine learning setting compares the predictive entropy against the expected entropy. Entropy is calculated as follows,

$$H(p) = -\sum_{k=1}^K p_k \log(p_k) \quad (4)$$

where p is the softmaxed network output, followed by the Expected $\mathbb{E}(p) = \hat{p}$. This allows us to measure the similarity between two random variables; if samples are very similar, we will get a high MI which is shown below in Eq. 5

$$\text{MI} = H(\hat{p}) - \mathbb{E}\left[\sum_{i=1}^K H(p_i)\right] \quad (5)$$

where K is the number of MC samples. It is also possible to use the Expected Kullback-Leibler Divergence (EKL):

$$\mathbb{E}[KL(\hat{p}||p)] = \mathbb{E}\left[\sum_{i=1}^K \hat{p}_i \log\left(\frac{\hat{p}_i}{p_i}\right)\right] \quad (6)$$

This term is similar to the MI but measures the expected divergence among the possible softmax outputs.

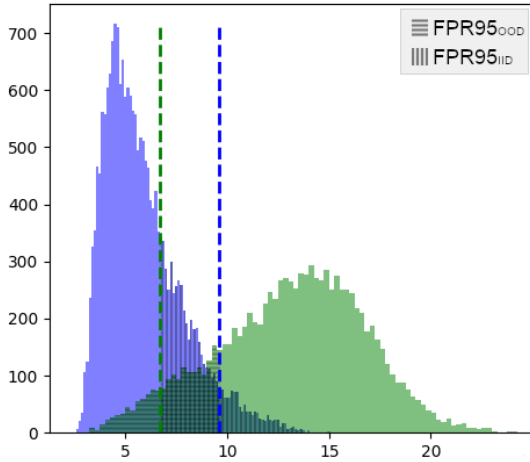


Fig. 3: Comparison of OOD and ID sample histograms with the green line at 5% quantile of ID data and the blue at 95% quantile of OOD data. With the y-axis shows the number of samples in that bin, and the x-axis is the value of the energy score.

Predictive variance is a more ad-hoc measure of uncertainty that evaluates the variance on the MC-sampled softmax outputs,

$$\sigma(p) = \mathbb{E} [(p - \hat{p})^2] \quad (7)$$

2) *Epistemic metrics*: The Area Under the Receiver Operating Characteristic curve (AUROC) depicts the relationship between the True Positive Rate (TPR) and the False Positive Rate (FPR). It can be interpreted as the probability that a positive sample is assigned a higher detection score than a negative example. The AUROC score is not affected by class imbalance, which is desirable.

The Area Under the Precision-Recall curve (AUPRC), on the other hand, has a common axis, the True Positive Rate (also known as Recall) but instead maps the relationship between that and the Precision (accuracy of the model). We show this metric with respect to both the ID and OOD datasets as the positive class (AUPRC_{ID/OOD}).

Finally, to waive any confusion in previous works about the use of FPR95, we clearly define FPR95_{ID} and FPR95_{OOD}. On the one hand, FPR95_{ID} is used when the positive samples are the ID dataset, i.e., the number of false positive cases when 95% of our data is correctly classified. On the other hand, the FPR95_{OOD} is used when the out-of-distribution dataset is set as the positive class. When 95% of positive samples are correctly classified, how many ID samples are within that range, this is visualized in figure 3.

B. FPR95 discussion

The FPR-N metric has been utilized in numerous contemporary studies on outlier detection. In certain instances, positive detection of outlier samples is demonstrated as FPR95_{OOD} [10], [29], [30], [31], which formally introduces FPR95 as an out-of-distribution measure. However, other studies, such as those by Liu et al. [32], [11], [12],

[28], employ the in-distribution class as the positive category (FPR95_{ID}). Finally, Liang et al. [14] and Hsu et al. [15] use True Negative Rate at 95% True Positive Rate (TNR@TPR95). These studies refer to one another but interchange the two methodologies for computing FPR without explicit clarification.

For FPR95_{OOD}, what is being tested is the proportion of ID samples with lower confidence when 95% of the outliers are found. This is, in essence, how good we are at detecting false alarms. The FPR95_{ID} metric shows us how big a proportion of the outliers are present when 95% of the ID data is correctly classified, resulting in a metric that describes how well we find ID samples. This choice is often vague, and while the metrics are similar, it highlights different things. That is how well we predict our in-distribution samples and how good we are at finding outliers. This also applies to the AUPR metric, where we have a big difference in what is shown depending on the positive selection.

V. EXPERIMENTAL SETUP

A. Backbone Network Structure

Our backbone model follows the convention set by [10] that is, the Wide ResNet (WRN) [33]; this architecture strikes a balance between performance and sensitivity. The Networks are trained for 100 epochs with a cosine annealing learning rate scheduler starting at 0.1 to produce models comparable to the VOS baseline. The Baseline WRN model achieves a 94.5% accuracy on CIFAR10 and 97.58% on our ships dataset. We follow the selected hyper-parameters to compare our method to the model trained in [28]. Thus the models with Logit Normalization are trained for 200 epochs with an initial learning rate of 0.1 with a step-wise scheduler reducing the learning rate by a factor of ten at 80 and 140 epochs. Both training setups have a batch size of 128 and are optimized using SGD with a momentum of 0.9 and weight decay of $5e - 4$. When MC sampling for inference, we use a dropout with a 10% chance of dropping a neuron.

B. Datasets

In this section, we describe the datasets used for our experiments, split into an 80/20 split of training/testing samples.

1) *CIFAR10*: For most of our experiments, we use the CIFAR10 dataset [34] as our ID dataset. The CIFAR10 dataset consists of 60,000 32x32 color images of 10 classes: airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. We use this simple but popular dataset as a baseline to compare our work against previous works.

2) *OOD Data*: To test our models against outliers, we need samples that are OOD. The authors of “ODIN” [14] proposed using a set of 6 outlier datasets. We will also follow the same practice commonly followed in the related literature. Therefore, as in [14], our six test datasets are: The Textures dataset, introduced by Cimpoi et al. [35], contains describable textural images. The SVHN dataset, proposed by Netzer et al. [36], comprises 32x32 color images of house numbers, with ten classes representing the digits 0-9. Zhou



TABLE I: Dataset description and example images of the SHIPS dataset

Annotation	Count T / V
Red buoy	789 / 194
Green buoy	566 / 113
Human / Kayak	771 / 170
Commercial vessel	3,020 / 947
Leisure craft	1,915 / 476
Sailboat (sail down)	1,539 / 358
Sailboat (sail up)	2,389 / 778

et al. [37] introduced the Places365 dataset, which comprises images for scene recognition rather than object recognition. The LSUN dataset, another scene understanding dataset, has fewer classes than Places365, and the cropped and resized versions are denoted as LSUN-C and LSUN-R, respectively. Finally, iSUN is a large-scale eye-tracking dataset selected from natural scene images in the SUN database [38] and was introduced by Xu et al. [39].

3) *SHIPS*: This work includes a self-collected and annotated dataset. This dataset is a finely classified version of our in-house datasets and data from other sources. Stets et al. in [40] introduced our in-house maritime dataset, which consists of 51,000 images with 31,900 annotations separated into two classes; boats and buoys. Of these, a subset has been selected and more granularly labeled.

We include data from other relevant datasets, such as the Singapore Maritime Dataset [41]. This dataset contains mixed samples of buoys and boats collected in Singaporean waters. Furthermore, we add data from the target area collected from online sources, primarily from videos and photos from Limfjorden (our area of interest). We add this to relevant data collected from publicly available data sets such as COCO, VOC, CIFAR10, and ImageNet.

We refer to this dataset as the **SHIPS** dataset and is made publicly available¹. The SHIPS dataset is a classification dataset consisting of seven unique classes. Refer to Table I for a description and examples of images of the dataset.

VI. RESULTS

This section shows the results gathered from testing the proposed method. Our initial results will primarily focus on our models trained on CIFAR10 to be more comparable with previous work. However, our main goal is to provide the method that has the best performance on our SHIPS dataset, which is not as curated as CIFAR10.

In Table III, we present the improvement that VOS and MC-Dropout yield on OOD detection. We note that the MC-Sampling allows us to regulate better incorrect outlier detections that VOS provides. As we increase the number of samples for MC-Dropout, we decrease the overall FPR95; this is shown to work with diminishing returns (we found the best performance to speed was 10 MC-Samples). Furthermore, the output, when predicting with MC-Samples provides a

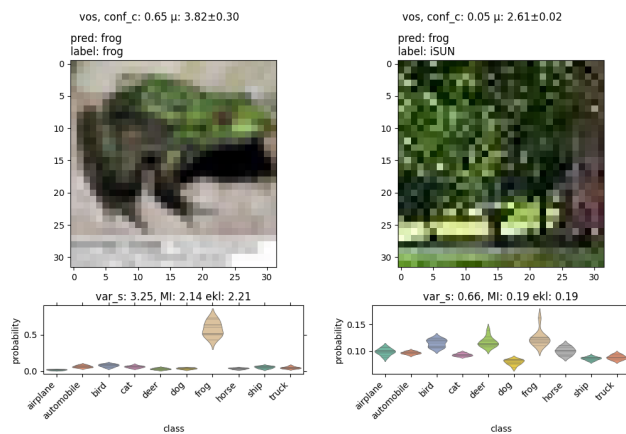


Fig. 4: On the left, we have an ID sample with a high MI but low energy. On the right is an OOD sample with Low MI and low Energy.

prediction bound displaying the variability of the estimates see Figure 4. This is also noticeable in the energy term, where samples with significant energy variations are samples with very low confidence. Ideally, we want our samples with high confidence to have a high energy score while also having a low mutual information score. From table IV, we show that the MI for incorrectly classified ID samples is higher than the correct counterpart, per our expectations. We use the pre-trained model presented in [12] as our baseline. However, we instead sample their model with MC-Dropout by enforcing dropout at inference. This shows a slight improvement with five samples; while sampling ten times, we achieve a score of 5.0% / 6.4% better than the baseline VOS. Our results indicate that by using MC-sampling on a VOS model, we allow for better filtering of the incorrect OOD samples, thus achieving better $FPR95_{ID/OOD}$ and AUROC/AUPR. We see a less impressive improvement for our SHIPS dataset, with a respective improvement of 41.0% / 10.1% for ID and OOD going from the baseline to VOS, and with MC-Dropout, we get a 44.2% / 20.1% improvement.

After introducing the Log normalization loss [28], we observed a significant decrease in the $FPR95_{ID/OOD}$ compared to the baseline VOS, from 24.87 to 11.91 (50.3%) and from 27.71 to 12.00 (57.0%). The performance of MC-Dropout improves with increasing sample size, but with diminishing returns. However, we noticed a different outcome when applying the same testing scheme to our SHIPS dataset. The LN substantially reduced the effectiveness of OOD detection while still maintaining good accuracy and being well-calibrated. We will delve deeper into this phenomenon in the discussion.

In table IV we show the relative difference in MI between the correctly classified and incorrectly classified samples and the relative difference between the IID and OOD datasets. We can see that the inclusion of Logit Normalization decreases the logit variation in the ID datasets and the OOD ID datasets by roughly 90%, an order of magnitude lower. We see this to a lesser extent on the SHIPS dataset (73%). The MI

¹<https://github.com/DTU-PAS/Ships-Classification>

TABLE II: Detailed performance of the baseline vs. our method

$D_{CIFAR10}$	baseline VOS / MC ₁₀ -LN-VOS					
	FPR95 _{ID}		AUPR _{ID}	AUROC	FPR95 _{OOD}	
<i>Texture</i>	46.32 ± 0.80 / 24.10 ± 0.76	95.93 / 98.44	86.60 / 94.12	66.91 ± 1.52 / 31.52 ± 2.06	64.99 / 82.25	
<i>Places365</i>	38.18 ± 1.19 / 32.39 ± 0.59	97.12 / 98.45	96.80 / 99.91	51.95 ± 3.70 / 25.33 ± 0.86	71.55 / 75.99	
<i>LSUN_C</i>	3.01 ± 0.34 / 0.13 ± 0.07	99.85 / 99.57	95.93 / 98.58	3.40 ± 0.24 / 0.20 ± 0.02	96.91 / 99.98	
<i>SVHN</i>	24.85 ± 0.47 / 7.20 ± 0.45	99.00 / 99.70	95.93 / 98.58	19.23 ± 0.58 / 6.90 ± 0.40	80.73 / 93.88	
<i>iSUN</i>	17.23 ± 0.78 / 3.30 ± 0.40	99.29 / 99.78	96.52 / 98.94	13.42 ± 0.72 / 3.89 ± 0.17	85.86 / 93.34	
<i>LSUN_R</i>	14.12 ± 0.73 / 3.66 ± 0.43	99.42 / 99.78	97.12 / 98.90	11.32 ± 0.47 / 4.13 ± 0.23	88.16 / 93.38	
<i>Average</i>	23.96 / 11.91	98.36 / 99.37	94.17 / 97.39	27.71 / 12.00	81.15 / 89.97	

TABLE III: Comparison of models on CIFAR10 and SHIPS

D_{ID}	CIFAR10				SHIPS			
	ID acc ↑	FPR95 ↓	AUROC ↑	AUPR ↑	ID acc ↑	FPR95 ↓	AUROC ↑	AUPR ↑
WRN	94.84	51.70 / 30.84	91.09	97.96 / 64.57	97.58	54.66 / 48.62	87.64	96.32 / 54.83
Baseline-VOS	94.62	23.96 / 27.71	93.86	96.37 / 80.62	97.58	32.26 / 44.17	90.45	95.29 / 74.80
MC ₅ -VOS	94.46	23.53 / 26.32	94.31	98.42 / 81.68	97.38	31.34 / 41.21	92.41	96.32 / 77.78
MC ₁₀ -VOS	94.46	22.91 / 25.91	95.09	98.13 / 83.32	97.32	30.64 / 39.41	92.41	96.32 / 77.98
LN	94.50	14.82 / 12.90	96.89	99.31 / 86.43	96.88	87.41 / 66.40	75.98	82.55 / 60.50
LN-VOS	94.60	12.54 / 12.13	97.32	99.36 / 89.54	96.68	88.99 / 68.90	74.12	80.65 / 59.50
MC ₁₀ -LN-VOS	94.65	11.91 / 12.00	97.39	99.37 / 89.97	96.68	76.27 / 62.50	79.91	84.45 / 69.52

TABLE IV: Scale difference of MI for the correctly classified ID samples versus the incorrectly classified ones $\frac{F}{T}(I)$. The whole ID dataset vs. the OOD datasets. We notice that the introduction of LN dramatically reduces the MI for OOD samples when training on CIFAR10, where it has the opposite effect for our SHIPS dataset $\frac{OOD}{ID}(II)$.

D_{ID}	CIFAR10		SHIPS	
	(I)	(II)	(I)	(II)
MC ₅ -VOS	11.80	4.27	7.23	4.44
MC ₁₀ -VOS	11.52	3.86	7.32	4.54
MC ₅ -LN-VOS	1.17	0.38	1.79	4.32
MC ₁₀ -LN-VOS	1.19	0.38	1.95	4.49

difference between the ID and OOD datasets for CIFAR10 is again roughly 90%, whereas, for the SHIPS dataset, we do not see any difference.

VII. DISCUSSION

The presented methods show a clear improvement in both outlier detection (epistemic uncertainty) and the improvement of intra-class confidence (aleatoric uncertainty). We want to explore issues with the current work further. We introduce new parameters that must be tuned, adding to the model complexity. MC-Sampling requires more computation as, in essence, the model is run multiple times. These issues are of interest, and a more unified method for parameter selection and alternate methods of more efficient MC-Sampling would need to be explored.

Although not our primary goal, we show that our method can achieve, to the best of our knowledge, results not yet achieved when no OOD data is available during training. The relationship between the MI of OOD and ID samples when LN is applied remains to be adequately explored. However, the application of the LN loss on the CIFAR10 dataset reduces the MI dramatically in OOD datasets see table IV. The combination of scaled MI and the energy

score, improved the performance of low-energy ID samples. Furthermore, by extending our method to include ODIN, we further improve our FPR95 results; 11.91 → 11.65 For ID and 12.00 → 10.67 OOD.

Our findings indicate a significant disparity between the performance of models trained on CIFAR10 and SHIPS datasets. We have observed that extending OOD models with MC-Dropout is a promising approach, and we aim to investigate its potential further. Our study highlights the crucial role that dataset characteristics play in determining the effectiveness of OOD detection methods. For instance, the SHIPS dataset is relatively small and skewed towards certain types of vessels, with a limited and high entropy representation of other objects, such as buoys and humans engaged in water-based activities (e.g., kayaking, rowing, swimming). Further investigation is needed to address these challenges, how to apply LN on poorly curated datasets, and to develop robust and reliable OOD detection methods that can perform well in diverse real-world scenarios.

VIII. CONCLUSION

This work has aimed to explore uncertainty estimation for the classification of maritime objects. We propose a method providing a holistic and usable uncertainty measure. We have presented our experimental setup, producing well-calibrated models while providing a usable aleatoric and epistemic uncertainty measure. Furthermore, we show that our method, when considering CIFAR10, to the best of our knowledge, performs 8% better than the previous state-of-the-art models only trained on ID data and 77% better than the vanilla WRN model. We show the varying performance of the applied methods when considering a well-curated dataset (CIFAR10) and a more applied dataset, SHIPS. Our proposed method performs well on the CIFAR10 and SHIPS datasets, performing 5% better than the baseline VOS, on both datasets and 55% / 44% better, respectively, than the vanilla WRN models.

REFERENCES

- [1] J. Beckett, F. Schöller, E. Boukas, M. Blanke, and L. Nalpantidis, "Lipschitz constrained neural networks for robust object detection at sea," *IOP Conference Series: Materials Science and Engineering*, vol. 929, no. 1, p. 012023, 11 2020.
- [2] J. Beckett, W. Seto, A. Deole, S. Bandyopadhyay, N. Rahimi, S. Talebi, M. Mesbah, and A. Rahmani, "Robust vision-based multi-spacecraft guidance navigation and control using CNN-based pose estimation," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 3 2022. [Online]. Available: <https://doi.org/10.1109/aero53065.2022.9843396>
- [3] J. Beckett, F. Schöller, E. Boukas, M. Blanke, and L. Nalpantidis, "Bolstering maritime object detection with synthetic data," *IFAC-PapersOnLine*, vol. 55, no. 31, pp. 64–69, 2022.
- [4] J. Beckett, E. Boukas, M. Blanke, and L. Nalpantidis, "Reweighting neural network examples for robust object detection at sea," *Electronics Letters*, vol. 57, no. 16, pp. 608–610, 4 2021.
- [5] J. Beckett, E. Boukas, and L. Nalpantidis, "Re-annotation of training samples for robust maritime object detection," *Machine Learning with Applications*, vol. 10, p. 100411, 12 2022.
- [6] C. M. Bishop, "Mixture density networks," 1994.
- [7] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *International Conference on Machine Learning*. PMLR, 2017, pp. 1321–1330.
- [8] B. Charpentier, D. Zügner, and S. Günemann, "Posterior network: Uncertainty estimation without ood samples via density-based pseudo-counts," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1356–1367, 2020.
- [9] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," *arXiv preprint arXiv:1610.02136*, 2016.
- [10] D. Hendrycks, M. Mazeika, and T. Dietterich, "Deep anomaly detection with outlier exposure," *arXiv preprint arXiv:1812.04606*, 2018.
- [11] W. Liu, X. Wang, J. Owens, and Y. Li, "Energy-based out-of-distribution detection," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 464–21 475, 2020.
- [12] X. Du, Z. Wang, M. Cai, and Y. Li, "VOS: Learning what you don't know by virtual outlier synthesis," *arXiv preprint arXiv:2202.01197*, 2022.
- [13] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," *Advances in neural information processing systems*, vol. 31, 2018.
- [14] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," *arXiv preprint arXiv:1706.02690*, 2017.
- [15] Y.-C. Hsu, Y. Shen, H. Jin, and Z. Kira, "Generalized ODIN: Detecting out-of-distribution image without learning from out-of-distribution data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [16] K. Lee, H. Lee, K. Lee, and J. Shin, "Training confidence-calibrated classifiers for detecting out-of-distribution samples," *arXiv preprint arXiv:1711.09325*, 2017.
- [17] S. Mohseni, M. Pitale, J. Yadawa, and Z. Wang, "Self-supervised learning for generalizable out-of-distribution detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5216–5223.
- [18] M. Grcić, P. Bevandić, and S. Šegvić, "Dense open-set recognition with synthetic outliers generated by Real NVP," *arXiv preprint arXiv:2011.11094*, 2020.
- [19] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 1321–1330. [Online]. Available: <https://proceedings.mlr.press/v70/guo17a.html>
- [20] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *International conference on machine learning*. PMLR, 2018, pp. 2796–2804.
- [21] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.
- [22] J. Lampinen and A. Vehtari, "Bayesian approach for neural networks—review and case studies," *Neural networks*, vol. 14, no. 3, pp. 257–274, 2001.
- [23] S. Kullback, *Information theory and statistics*. Courier Corporation, 1997.
- [24] L. Smith and Y. Gal, "Understanding measures of uncertainty for adversarial example detection," *arXiv preprint arXiv:1803.08533*, 2018.
- [25] G. Wang, W. Li, M. Aertsen, J. Depreest, S. Ourselin, and T. Vercauteren, "Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks," *Neurocomputing*, vol. 338, pp. 34–45, 2019.
- [26] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [27] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.
- [28] H. Wei, R. Xie, H. Cheng, L. Feng, B. An, and Y. Li, "Mitigating neural network overconfidence with logit normalization," *arXiv preprint arXiv:2205.09310*, 2022.
- [29] D. Hendrycks, M. Mazeika, S. Kadavath, and D. Song, "Using self-supervised learning can improve model robustness and uncertainty," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/a2b15837edac15df90721968986f7f8e-Paper.pdf>
- [30] Y. Li and N. Vasconcelos, "Background data resampling for outlier-aware classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [31] S. Thulasidasan, S. Thapa, S. Dhaubhadel, G. Chennupati, T. Bhat-tacharya, and J. Bilmes, "An effective baseline for robustness to distributional shift," in *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2021, pp. 278–285.
- [32] S. Liu, R. Garrepalli, T. Dietterich, A. Fern, and D. Hendrycks, "Open category detection with PAC guarantees," in *International Conference on Machine Learning*. PMLR, 2018, pp. 3169–3178.
- [33] S. Zagoruyko and N. Komodakis, "Wide residual networks," *CoRR*, vol. abs/1605.07146, 2016. [Online]. Available: <http://arxiv.org/abs/1605.07146>
- [34] A. Krizhevsky, G. Hinton, *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [35] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi, "Describing textures in the wild," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 3606–3613.
- [36] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," 2011.
- [37] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, "Places: A 10 million image database for scene recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 6, pp. 1452–1464, 2017.
- [38] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *2010 IEEE computer society conference on computer vision and pattern recognition*. IEEE, 2010, pp. 3485–3492.
- [39] P. Xu, K. A. Ehinger, Y. Zhang, A. Finkelstein, S. R. Kulkarni, and J. Xiao, "Turkergaze: Crowdsourcing saliency with webcam based eye tracking," *arXiv preprint arXiv:1504.06755*, 2015.
- [40] J. D. Stets, F. E. T. Schöller, M. K. Plenge-Feidenhans'l, R. H. Andersen, S. Hansen, and M. Blanke, "Comparing Spectral Bands for Object Detection at Sea using Convolutional Neural Networks," *Journal of Physics: Conference Series*, vol. 1357, p. 12036, oct 2019.
- [41] D. K. Prasad, D. Rajan, L. Rachmawati, E. Rajabally, and C. Quek, "Video processing from electro-optical sensors for object detection and tracking in a maritime environment: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 8, pp. 1993–2016, 2017.