

Received March 27, 2020, accepted April 19, 2020, date of publication May 4, 2020, date of current version May 15, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2990985

A Computational Framework for Iceberg and Ship Discrimination: Case Study on Kaggle Competition

XULEI YANG¹, (Senior Member, IEEE), AND JIE DING^{2,3}

¹Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR), Singapore 138632

²Adaptive Networks and Control Laboratory, Department of Electronic Engineering, Fudan University, Shanghai 200433, China

³Research Center of Smart Networks and Systems, School of Information Science and Engineering, Fudan University, Shanghai 200433, China

Corresponding author: Jie Ding (dingjie@fudan.edu.cn)

ABSTRACT Iceberg and ship identification in satellite synthetic aperture radar (SAR) data plays an important role in offering an operational iceberg surveillance program. Here, the identification aims to detect ocean SAR targets and then categorize these targets into iceberg, ship, or unknown. Although the adaptive threshold techniques have achieved promising results on the ship and iceberg detection in SAR images, the discrimination between these two target classes is still very challenging for operational scenarios. This study presents a computational framework for iceberg and ship discrimination based on an ensemble of various deep learning and machine learning algorithms. On one hand, latest deep neural networks – namely, DenseNet and ResNet – are deployed in this study for end-to-end feature extraction and image classification directly on original SAR images. On the other hand, handcrafted features are extracted on de-speckled SAR images, followed by classification using advanced machine learning algorithms – namely, XGBoost and LightGBM. The outcomes from both sides are then combined through min-max median stacking approach to classify the given SAR images into iceberg and ship categories. The proposed framework has recently been deployed as the key kernel for the “Statoil/C-CORE Iceberg Classifier Challenge” organized by Kaggle. The performance is promising as our final scores were ranked 26 and 39 out of 3343 teams on public and private leaderboards, respectively. We hope that by sharing the solutions, we can further promote research interests in the field of iceberg and ship identification.


INDEX TERMS Iceberg and ship identification, satellite synthetic aperture radar (SAR) images, image classification.

I. INTRODUCTION

For oceanographic observations, satellite synthetic aperture radar (SAR) plays an important role, as it is able to monitor the oceans and floating structures in all weather conditions by using its active radar. Drifting icebergs present a threat to navigation and activities in areas such as offshore of the East Coast of Canada. To offer an operational iceberg surveillance program, it is essential to utilize satellite SAR data to identify between iceberg and ship. However, this process could be labor intensive, subjective, and error prone because satellite SAR data with coarser resolution is not as intuitive as satellite optical data for manually interpreting target classification. Therefore, it is desired to develop an automated method [1] for iceberg or ship identification. The identification involves

the detection of ocean SAR targets and then the classification of these targets.

There is extensive literature on ship detection using SAR, and much of this has been summarized in [2] and [3]. These ship detection algorithms were applied for the detection of icebergs from SAR data. One of the simple but widely used methods is the threshold-based detection algorithm. The algorithm can be thought of as a bright target detection on the sea, such that a pixel is selected as long as it passes the criteria (or threshold), regardless of what it represents. This algorithm can achieve success in ocean target detection because ships and icebergs usually appear as bright targets against the dark background of the ocean. However, it is not able to classify the target, and hence discrimination algorithms have to be further employed to label the target as a ship, an iceberg or unknown.

The associate editor coordinating the review of this manuscript and approving it for publication was Byung-Gyu Kim .

The previous methods on SAR images classification include using multivariate approaches, support vector machines (SVMs), and convolutional neural networks (CNNs). The combination of HV (transmit horizontally and receive vertically) and HH (transmit/receive horizontally) bands, can be used for feature extraction as shown in [4], from which it is discovered that ships are more likely to generate reflections on the edges in the HH bands. To distinguish between ships and icebergs, the discrimination between the two classes is carried out after edge detection by extracting features and conducting target classification. The model in [5] uses the Bayes rule to maximize the posteriori probability regarding two classes to achieve an accuracy of 93% for classifying ships. To discriminate ships from icebergs in simulated, dual polarized, medium resolution SAR data, a SVM classifier is proposed in [6] based on the features of intensity and polarimetric parameters. In [7], ships and icebergs are discriminated based on their different dominant scattering mechanisms.

In recent years, CNNs [8] have been leading the way in solving many challenging image classification problems. Due to the advent of AlexNet [9], which won 2012 ImageNet Large-Scale Visual Recognition Challenge, CNNs have been successfully employed in many image classification tasks [10], [11]. Compared to using hand-crafted features in conventional machine learning approaches, deep neural networks are able to automatically learn complex representations from input data. Recently, CNNs have achieved promising performance in SAR classification tasks [12], [13].

Specifically, for the task of ship and iceberg discrimination, Bentes *et al.* [14] applied CNNs to ship-iceberg discrimination and tested on TerraSAR-X StripMap images. Schwegmann *et al.* [15] employed a specific type of deep neural networks – the highway networks – to ship discrimination in SAR images and achieved promising results. Qdegaard *et al.* [16] used CNNs to detect ships in harbor background in SAR images. To circumvent the lack of training samples, they utilized a simulation software to generate simulated data for training. Song *et al.* [17] followed this idea and introduced a deep generative neural network for SAR automatic target recognition. Zhang *et al.* [18] proposed a complex-valued (CV) CNNs specifically designed to process complex values in PolSAR data, i.e., the off-diagonal elements of coherency or covariance matrix.

In this paper, a computational framework is proposed for iceberg and ship discrimination based on an ensemble of various deep learning and machine learning algorithms. On one hand, latest deep neural networks – namely, ResNet [19] and DenseNet [20] – are deployed in this study for end-to-end feature extraction and image classification directly on original SAR images. On the other hand, handcrafted features are extracted on de-speckled SAR images, followed by classification using advanced machine learning algorithms – namely, XGBoost [21] and LightGBM [22]. The outcomes from both approaches are then combined through min-max

median stacking approach to classify the given SAR images into iceberg and ship.

Generally, ensemble of models is a standard approach in applied machine learning to ensure that the most stable and best possible prediction is made. For this specific task, we found out that the deep features extracted by CNNs did not dominate the handcrafted features used in boosting methods. In other words, CNNs and boosting methods may discover diverse meaningful features, leading to different predictions for given SAR images. This is the main motivation of the combination of the predictions from CNNs and boosting methods. We make an assumption in this study that “CNNs and boosting methods may be good in different ways, and make different prediction errors”, as long as the majority of the models make correct predictions, the combination will lead to a better result. This is in consonance with the observation in [23]: “the reason that model ensemble works is that different models will usually not make all the same errors on the test set”.

The main contributions of this paper are three folds: Firstly, we propose a computational framework for iceberg and ship classification through the ensemble of advanced deep learning and machine learning techniques. Secondly, we address several unique challenges involved in target discrimination from coarse resolution satellite SAR images, such as data augmentation and pseudo labelling. Lastly, we publish the source codes¹ for public sharing as part of the Kaggle competition “Statoil/C-CORE Iceberg Classifier Challenge” [24]. Interested readers may re-use the source codes for their own target discrimination from SAR images.

The rest of this paper is organized as follows: The proposed iceberg and ship discrimination framework is presented in Section II. Case study of our proposed framework on Kaggle competition for ship and iceberg classification is illustrated in Section III. Insights and discussions related to target discrimination from SAR images are included in Section IV. Lastly, the conclusion is given in Section V.

II. METHODOLOGY

The goal of this paper is to develop automated ways to discriminate icebergs from ships for a given SAR image. The overall workflow of the proposed framework is illustrated in Fig.1, which consists of image classification by deep learning approaches (described in Section II-A), image de-speckling followed by machine learning classification (described in Section II-B), and the ensemble of outcomes from both deep learning and machine learning (described in Section II-C).

A. DEEP LEARNING APPROACHES

In recent years, the community has witnessed an accelerated growth in the development of deep learning techniques [25], [26] especially the deep CNNs like ResNet [19] and DenseNet [20]. These developments have significantly

¹ https://bitbucket.org/xulei_yang/iceberg_ship/src

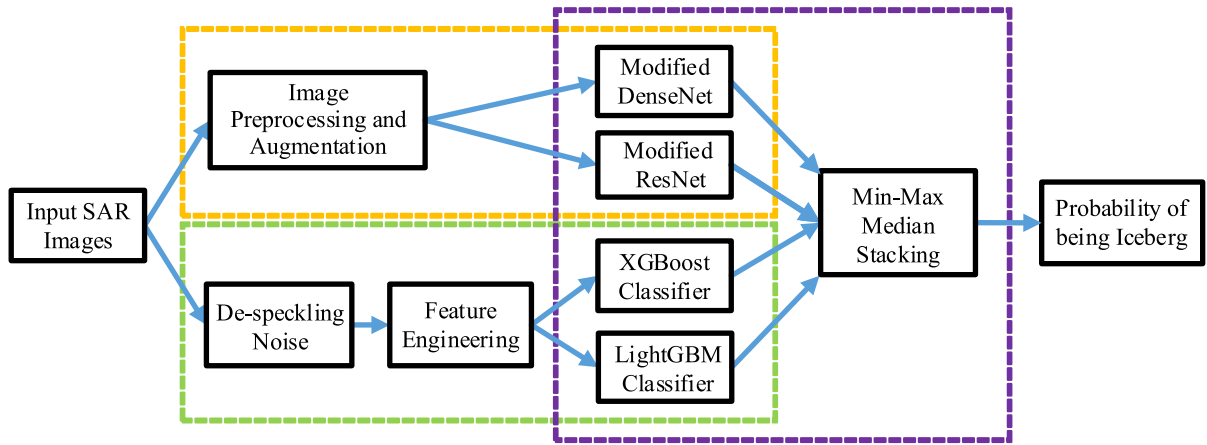


FIGURE 1. Overall workflow of the proposed computational framework.

pushed the frontier of the state-of-art in many application domains, including visual object detection, speech recognition, etc.

When it comes to neural network design, the trend in the past few years has pointed in one direction: deeper. However, deep networks are hard to train because of the notorious vanishing gradient problem where the gradient becomes infinitesimally small through repeated multiplication in the back-propagation process. As a result, as the network goes deeper, its performance gets saturated or even starts degrading rapidly [27].

ResNet [19] uses a so-called “identity shortcut connection” that skips one or more layers. In addition, it refines the residual block and proposed a pre-activation variant of the residual block such that the gradients can flow through the shortcut connections to any other earlier layer unimpededly. In such a way, ResNet makes it possible to train up to hundreds or even thousands of layers with compelling efficiency [28].

Huang *et al.* [20] proposed a novel architecture called DenseNet that further exploits the effects of shortcut connections - it connects all layers directly with each other. In this novel architecture, the input of each layer consists of the feature maps of all earlier layers, and its output is passed to each subsequent layers. The feature maps are aggregated with depth-concatenation. Other than tackling the vanishing gradients problem, this architecture also encourages feature reuse, making the network highly parameter-efficient.

In this subsection, we describe how we modified the ResNet and DenseNet models for iceberg and ship discrimination. Our approach is illustrated in Fig.2. The convolutional layers from ResNet and DenseNet are used to extract various layers of features from SAR images, which are then combined together for image classification. We use pre-trained ResNet and DenseNet models on ImageNet dataset. The pre-trained Deep CNN (DCNN) models require 3-channel input. As the original SAR images from Kaggle competition only contain two channels (HH and HV), the average of these two channels is used as the third channel. All the fully-connected layers in

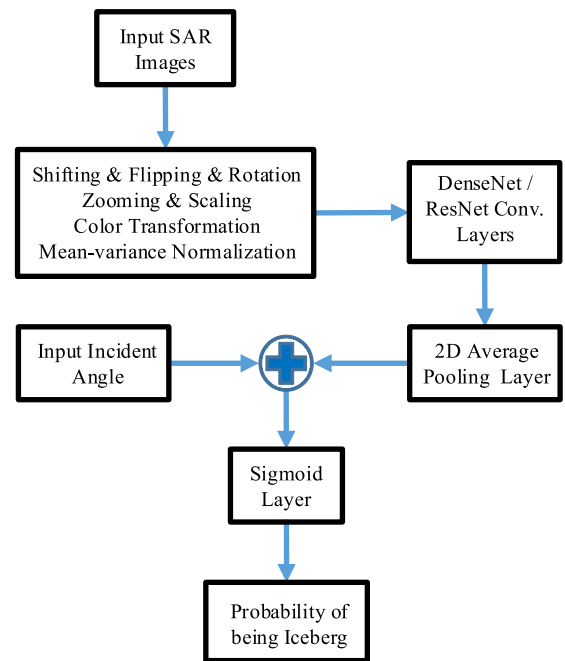


FIGURE 2. Diagram of deep learning approaches for ship and iceberg classification.

the original models are removed, and a 2D average pooling layer is added after the last convolutional layer, followed by a flatten layer before the last dense layer. Finally, Sigmoid activation is used in the last dense layer instead of the original Softmax function to obtain the probability of the image belonging to the iceberg class. In the Kaggle competition, incident angle of each SAR image is also provided. To make use of this information, the angle value is concatenated into the network just before the last dense layer, as shown in Fig.2.

B. MACHINE LEARNING METHODS

The workflow of using machine learning methods for ship and iceberg discrimination is illustrated in Fig. 3, which involves noise de-speckling, feature engineering, and image classification. The details are discussed in the ensuing text.

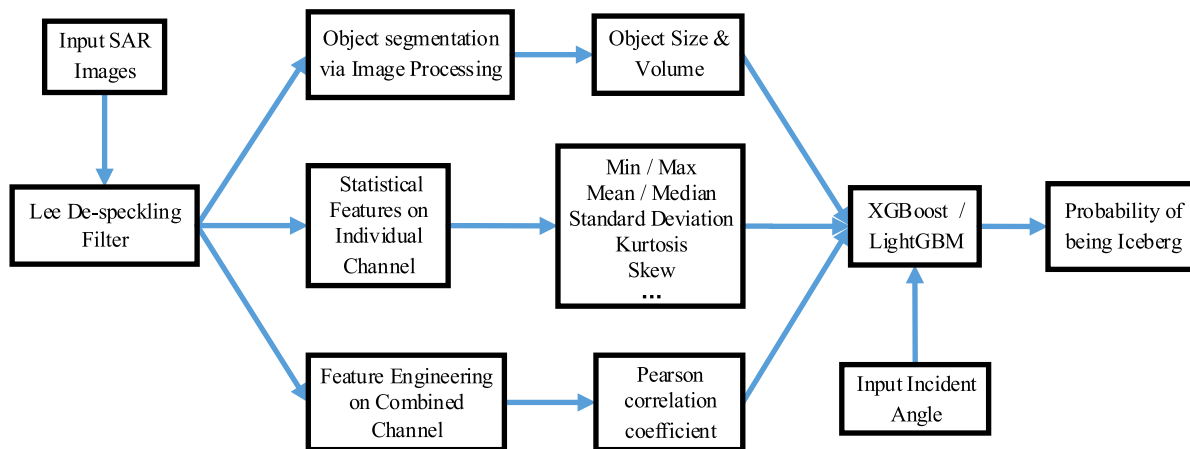


FIGURE 3. Diagram of machine learning approaches for ship and iceberg classification.

1) SAR IMAGES

SAR utilizes microwave-frequency light to actively image surface features [29]. One of the consequences is that the features being imaged often have a roughness with length scales similar to the wavelength of the light being used. So, the backscattered light will experience mutual interference creating speckle noise. Speckle noise presents itself as patches that are lighter or darker than they otherwise would have been based on the features being imaged, as demonstrated by Fig. 4 and Fig. 5.




FIGURE 4. Demonstration of filtering out speckle noise on a ship image before filtering. Left: ship band 1; right: ship band 2.




FIGURE 5. Demonstration of filtering out speckle noise on a ship image after filtering. Left: ship band 1; right: ship band 2.

Unlike deep learning approaches, which provide an end-to-end solution for the iceberg and ship discrimination

directly on original SAR images, machine learning classification methods require feature engineering to extract meaningful features to be used for classification. Raw SAR images from Kaggle competition contain multiple speckle noise (as shown in Fig. 4 and Fig. 5) which may affect feature extraction and image classification. In this work, the additive noise lee de-speckling filter [30] is used to reduce speckle noise before employing machine learning approaches.

2) DE-SPECKLING

A number of de-speckling methods, such as [31] and [32], have been presented to filter out speckle noise in SAR images. In this study, we employ the speckle noise reduction method called Lee Filter [30] to reduce such noise. It is assumed that speckle noise is additive with a constant mean of zero, a constant variance, and drawn from a Gaussian distribution. Using a window ($I \times J$ pixels) to scan the image with a stride of 1 pixel, the de-speckled value of the pixel in the center of the window located in the i th row and j th column is

$$z_{ij}^* = u_k + w * (z_{ij} - u_k) \tag{1}$$

where u_k is the mean value of all pixels in the window centered on pixel (i, j) , z_{ij} is the unfiltered value of the pixel, and w is a weight calculated by

$$w = \frac{\text{var}_k}{\text{var}_k + \text{var}_{\text{noise}}} \tag{2}$$

where var_k is the variance of all pixels in the window and $\text{var}_{\text{noise}}$ is the variance of the speckle noise. A possible alternative to using the actual value of the center pixel for z_{ij} is to use the median pixel value in the window. The parameters of the filter are the window/kernel size and the variance of the noise (which is unknown but perhaps can be estimated from the image as the variance over a uniform feature smooth like the surface of still water). Using a larger window size and noise variance will increase radiometric resolution at the expense of spatial resolution. As demonstrated by two images in Fig.5, most of the speckle noise in the SAR images are

successfully removed. We will perform feature extraction and image classification based on filtered images.

3) FEATURE ENGINEERING

The following features are extracted and used for this study:

- Regular aggregations per band considered as a signal:
 - Minimum
 - Maximum
 - Mean
 - Standard deviation
 - Kurtosis
 - Skew
- Other aggregations per band considered as an image:
 - Standard deviation after Sobel filtering on x
 - Standard deviation after Sobel filtering on y
 - Standard deviation after Laplace filtering
- Features from combined bands:
 - Pearson correlation coefficient
 - Standard deviation of $\sqrt{(\text{band } 1^2 + \text{band } 2^2)}$

4) IMAGE CLASSIFICATION

Tree boosting is a highly effective and widely used machine learning method. In this study, we employ the latest developments of boosting machines, i.e., XGBoost [21] and LightGBM [22] for the discrimination of ship and iceberg. XGBoost [21] stands for eXtreme Gradient Boosting, an algorithm that has recently been dominating applied machine learning and Kaggle competitions for structured or tabular data. XGBoost is an implementation of gradient boosted decision trees designed for speed and performance, which scales beyond billions of examples using far fewer resources than existing systems. LightGBM [22] is a gradient boosting framework that uses tree based learning algorithm. LightGBM grows tree vertically while other algorithm grows trees horizontally meaning that LightGBM grows tree leaf-wise while other algorithm grows level-wise. It will choose the leaf with max delta loss to grow. When growing the same leaf, leaf-wise algorithm can reduce more loss than a level-wise algorithm, hence resulting in much better accuracy which would otherwise be rarely achievable by any of the existing boosting algorithms. As shown in Fig. 3, by using the extracted features as described in the last subsection, together with the input feature of incident angle, XGBoost and LightGBM classifiers are used to classify the given SAR image into ship and iceberg by outputting the probability of the image being an iceberg.

C. STACKING

Different stacking methods have been used to build an ensemble from multiple classifiers. For this specific application, it is found that the simple min-max median staking generates the best performance. Assume that we have N outcomes from N classifiers and each classifier generates a probability value between $[0, 1]$ for a given SAR image. High probability indicates the image has high potential of being an iceberg, while low probability implies high potential of being a ship.

The basic idea behind min-max median is to select very high probability (larger than pre-defined max value, most likely being an iceberg) and very low probability (smaller than pre-defined min value, most likely being a ship) values first, then take median of these selected values. If no values are selected, then take median through all the N probability values. Effectiveness of the min-max median stacking will be demonstrated by the experiments in the next section.

III. EXPERIMENTS

A. DATASET

The Centre for Cold Ocean Resource Engineering (C-CORE) and Norwegian energy company Statoil have launched a competition to find a more effective method of spotting icebergs that pose a risk to ships and infrastructure. The ‘‘Iceberg Classifier Challenge’’ requires participants to find an ‘‘algorithm that automatically identifies if a remotely sensed target is a ship or iceberg,’’ according to the competition’s home page [24]. The challenge was posted on the Google-owned website Kaggle. Kaggle is used for hosting competitions where participants are invited to devise solutions to a wide range of data-related problems.

In this binary classification task, a participant has access to readings of radar backscatter to distinguish ships and icebergs. The data is of 2-channel, each polarized to transmit horizontally-receive horizontally/vertically (HH/HV). The readings are floating point values on the dB scale. Also included is the angle of incidence for each reading. There are 1604 samples in the training set and 8424 samples in the testing set. For testing set, 20% of the samples are used for public leaderboard scoring and the rest for private leaderboard scoring. The classes in the training set are fairly well-balanced wherein 900 are ships and the rest are icebergs. Fig.6 and Fig. 7 show some randomly selected samples from the training set for 2 ships and 2 icebergs, respectively, where pseudo color is used for visualization. It can be seen from the figures that it could be hard to discern some ships from icebergs by visual inspection of a channel that had been coerced into pseudo-color images.

B. EVALUATION METRICS

Submissions are evaluated on the log loss between the predicted values and ground truth. For Kaggle competition, each image has been labeled with one true class (not disclosed for testing set). For scoring, we must submit a set of predicted probabilities (one for each image). The performance is then evaluated by the log loss function as follows:

$$\log \text{ loss} = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m y_{ij} \log(p_{ij}) \quad (3)$$

where n is the number of images in the test set, m is the number of image class labels, \log is the natural logarithm, y_{ij} is 1 if observation i belongs to Class j and 0 otherwise, and p_{ij} is the predicted probability that observation i belongs to Class j . More specifically, for two-class image classification




FIGURE 6. Image samples of ships from Kaggle competition. (a) Ship sample 1. Left: band 1; right: band 2. (b) Ship sample 2. Left: band 1; right: band 2.




FIGURE 7. Image samples of icebergs from Kaggle competition. (a) Iceberg sample 1. Left: band 1; right: band 2. (b) Iceberg sample 2. Left: band 1; right: band 2.

problem, the log loss function is as follows:

$$\log loss = \frac{1}{n} \sum_{i=1}^n [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (4)$$

where y_i is 1 if the image is an iceberg, and 0 if it is a ship. The variable \hat{y}_i is the predicted probability of the image being an iceberg. A smaller log loss is better.

C. TRAINING

The deep learning models were constructed on the Keras deep learning platform with Tensorflow backend. The Adam optimizer with learning rate $\eta = 0.0001$, $\beta_1 = 0.9$, and $\beta_2 = 0.999$ was implemented for compiling the Keras model. The value of fuzz vector used in numeric expressions was set as $1e - 08$. All weights were initialized with 0 mean and 0.1 standard deviation Gaussian distribution. These weights are optimized by back propagation through the aforementioned loss functions to calculate the penalty between the prediction and the ground truth in every batch. The network was trained for 50 epochs with a mini batch size of 32 using a 5-fold cross validation. The final accuracy is estimated by averaging 5 different values produced by each fold. All the experiments were conducted in Ubuntu 16.04 with 4 Nvidia GTX1080 GPU cards.

D. RESULTS

We tested each of the discussed methods on Kaggle iceberg classification dataset [24]. The generated score values are listed in Table 1. Overall, the deep learning approaches performed much better than machine learning methods, e.g., ResNet achieved log loss 0.1494 and 0.1474, much smaller than 0.1882 and 0.2063 by XGBoost, on public and private leaderboard, respectively. By combining outcomes from both deep learning and machine learning, our proposed framework achieved the best performance, i.e., log loss values of 0.1070 and 0.1295, which ranked 26 and 39 out of 3343 international teams on public and private leaderboard, respectively.

IV. DISCUSSION

Ship and iceberg discrimination from low resolution SAR images is a challenging task. In this section, we discuss several interesting findings from our participation in the competition as follows. Several of them have also been deployed in our previous kaggle competitions [33], in which a two-stage detection scheme was proposed to handle small object recognition from large background within a given image.

A. DATA AUGMENTATION IS ESSENTIAL FOR CNNs CLASSIFIERS

As shown in the experiment results, image augmentation is one of the most important methods which can help boost the performance of deep neural networks with limited training data. To build good image classifiers, different data augmentation strategies can be adopted for different domain applications. In this work, a combination of image processing techniques, such as random rotation, shifts, shear, scale and flips, was leveraged to create artificial images. Meanwhile, mean-variance normalization, as well as color space transformation and elastic transformation were also implemented to

TABLE 1. Performance of each method on Kaggle iceberg classification dataset [24].

Methods	Public Leaderboard	Private Leaderboard	Remarks
XGBoost	0.1882	0.2063	Ensemble of various XGBoost models
LightGBM	0.1807	0.2088	Ensemble of various LightGBM models
ResNet	0.1494	0.1474	Ensemble of various ResNet18 & ResNet50 models
DenseNet	0.1408	0.1449	Ensemble of various DenseNet121 models
Min-max stacking	Median 0.1070	0.1295	

enhance the image augmentation. For more details on image augmentation, the reader may refer to [34].

B. CROSS-VALIDATION HELPS AVOID MODEL OVERFITTING

The goal of a deep learning or machine learning model is to generalize the training data to any data different from problem domain. In this Kaggle challenge, the number of samples in the training dataset is much smaller than that in the testing dataset. As a result, the trained model may yield higher accuracy on the training set than the testing dataset. In other words, there is a problem with overfitting. To mitigate this problem, we employed k -folder ($k = 5$) cross-validation wherein the output of the model is averaged over k -trained sub-models. The experiments showed that this model generalization strategy increased our competition ranking to some extent in both public and private leaderboards.

C. PSEUDO LABELING INCREASES TRAINING SAMPLES

In many machine learning domains, data acquisition is one of the most expensive steps since many time-consuming and labor-intensive processes are involved. For example, ground truth masks have to be drawn manually and further confirmed by many other experts through voting. Therefore, using limited data and yet able to train a good model is still an open problem. In this work, the pseudo labeling approach was applied where the unlabeled data in the testing dataset were utilized for the training process. Specifically, a model is first trained based on the training dataset. Then we use such a model to perform the prediction task on the testing dataset where the predicted probability for each testing sample is calculated. After that, the most confident samples from the testing dataset are in turn added into the training dataset. This process can be iterated to help enlarge the training dataset progressively. It is shown that the performance of the trained model was improved by using pseudo labeling from the experiments.

D. DE-SPECKLING WORKS FOR MACHINE LEARNING ONLY

It is interesting to find out that deep learning obtains worse performance on de-speckled data samples as compared to raw SAR images. We believe this is due to the fact that

the de-noising process may discard some useful structures around the ship or iceberg objects, which could otherwise be learned by the deep neural networks to assist the classification of ship and iceberg. As such, we train ResNet and DenseNet models directly using raw SAR images. For machine learning methods, hand-crafted features could be obfuscated by speckle noise. Hence, it is difficult to find meaningful features directly on noisy raw SAR images as compared to de-speckled images. This explains why de-speckling improves the performance of XGBoost and LightGBM classifiers on ship and iceberg classification.

E. ENSEMBLE BOOSTS PERFORMANCE

In the past years, as demonstrated in a number of computer vision competitions, deep learning approaches have achieved superior performance over traditional machine learning approaches on most of the image and video analysis tasks. It is natural that a deep learning option would be of first priority for this image-related competition. Machine learning approaches, while performing worse than CNNs on image classification tasks (as shown in Table 1), may serve to add meaningful features through delicate feature engineering to assist the classification of deep learning. As shown in the experimental results, the combination of deep learning and machine learning leads to a better performance as compared to that of individual models.

V. CONCLUSION

In this study, we presented a computational framework to tackle the ship and iceberg discrimination problem by combining the advantages of both deep learning and machine learning techniques. We also addressed several unique practical challenges related to this specific real-world application. The proposed approach was evaluated through the case study on the recent Kaggle competition for ship and iceberg classification. Our results were ranked nearly top 1% on both public and private leaderboards, which verified the effectiveness of the proposed approach.


REFERENCES

- [1] A. K. Mazur, A. K. Wählin, and A. Krężel, "An object-based SAR image iceberg detection algorithm applied to the amundsen sea," *Remote Sens. Environ.*, vol. 189, pp. 67–83, Feb. 2017.
- [2] D. J. Crisp, "The state-of-the-art in ship detection in synthetic aperture radar imagery," Defence Sci. Technol. Org. Salisbury (Australia) Info Sci. Lab, Salisbury, SA, Australia, Tech. Rep. DSTO-RR-0272, 2004.

- [3] A. Marino, M. Sanjuan-Ferrer, I. Hajnsek, and K. Ouchi, "Ship detection with spectral analysis of synthetic aperture radar: A comparison of new and well-known algorithms," *Remote Sens.*, vol. 7, no. 5, pp. 5416–5439, 2015.
- [4] M. Modava and G. Akbarizadeh, "A level set based method for coastline detection of SAR images," in *Proc. 3rd Int. Conf. Pattern Recognit. Image Anal. (IPRIA)*, Apr. 2017, pp. 253–257.
- [5] C. Howell, J. Mills, D. Power, J. Youden, K. Dodge, C. Randell, S. Churchill, and D. Flett, "A multivariate approach to iceberg and ship classification in HH/HV ASAR data," in *Proc. IEEE Int. Symp. Geosci. Remote Sens.*, Jul. 2006, pp. 3583–3586.
- [6] M. Denbina, M. J. Collins, and G. Atteia, "On the detection and discrimination of ships and icebergs using simulated dual-polarized RADARSAT constellation data," *Can. J. Remote Sens.*, vol. 41, no. 5, pp. 363–379, Sep. 2015.
- [7] C. Howell, J. Youden, K. Lane, D. Power, C. Randell, and D. Flett, "Iceberg and ship discrimination with ENVISAT multi-polarization ASAR," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, 2004, pp. 113–116.
- [8] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [9] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, May 2017.
- [10] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [11] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1–9.
- [12] X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, Dec. 2017.
- [13] L. Zhang, L. Zhang, and B. Du, "Deep learning for remote sensing data: A technical tutorial on the state of the art," *IEEE Geosci. Remote Sens. Mag.*, vol. 4, no. 2, pp. 22–40, Jun. 2016.
- [14] C. Bentes, A. Frost, D. Velotto, and B. Tings, "Ship-iceberg discrimination with convolutional neural networks in high resolution sar images," in *Proc. 11th Eur. Conf. Synth. Aperture Radar (EUSAR)*, 2016, pp. 1–4.
- [15] C. P. Schwegmann, W. Kleynhans, B. P. Salmon, L. W. Mdakane, and R. G. V. Meyer, "Very deep learning for ship discrimination in synthetic aperture radar imagery," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2016, pp. 104–107.
- [16] N. Odegaard, A. O. Knapskog, C. Cochin, and J.-C. Louvigne, "Classification of ships using real and simulated data in a convolutional neural network," in *Proc. IEEE Radar Conf. (RadarConf)*, May 2016, pp. 1–6.
- [17] Q. Song, F. Xu, and Y.-Q. Jin, "Deep SAR image generative neural network and auto-construction of target feature space," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, 2017.
- [18] Z. Zhang, H. Wang, F. Xu, and Y.-Q. Jin, "Complex-valued convolutional neural network and its application in polarimetric SAR image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 12, pp. 7177–7188, Dec. 2017.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [20] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4700–4708.
- [21] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 785–794.
- [22] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "LightGBM: A highly efficient gradient boosting decision tree," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 3146–3154.
- [23] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>
- [24] (2017). *Kaggle Competition: Statoil/c-Core Iceberg Classifier Challenge*. [Online]. Available: <https://www.kaggle.com/c/statoil-iceberg-classifier-challenge>
- [25] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [26] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Netw.*, vol. 61, pp. 85–117, Jan. 2015.
- [27] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, 2010, pp. 249–256.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2016, pp. 630–645.
- [29] J.-M. Nicolas and F. Adragna, "The principles of synthetic aperture radar," in *Processing of Synthetic Aperture Radar Images*. London, U.K.: Wiley, 2008, pp. 25–55.
- [30] J.-S. Lee, "Speckle analysis and smoothing of synthetic aperture radar images," *Comput. Graph. Image Process.*, vol. 17, no. 1, pp. 24–32, Sep. 1981.
- [31] F. Argenti, A. Lapini, T. Bianchi, and L. Alparone, "A tutorial on speckle reduction in synthetic aperture radar images," *IEEE Geosci. Remote Sens. Mag.*, vol. 1, no. 3, pp. 6–35, Sep. 2013.
- [32] G. Chierchia, D. Cozzolino, G. Poggi, and L. Verdoliva, "SAR image despeckling through convolutional neural networks," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2017, pp. 5438–5441.
- [33] X. Yang, Z. Zeng, S. G. Teo, L. Wang, V. Chandrasekhar, and S. Hoi, "Deep learning for practical image recognition: Case study on Kaggle competitions," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 923–931.
- [34] (2018). *Imgaug—Augmenting Images for Machine Learning*. [Online]. Available: <https://github.com/aleju/imgaug>



XULEI YANG (Senior Member, IEEE) is currently a Senior Research Scientist and the Programme Head with the Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR), with more than 13 years of research and development experiences in image/signal analysis and deep/machine learning. He is also a Kaggle Competition Master. He is also the principal investigator for various projects involved in providing deep learning solutions for healthcare, Fintech, and computer vision. He has published more than 60 scientific articles and international patents. His current research interests focus on deep learning for image/signal analysis, recommendation, and prediction.



JIE DING received the B.Eng. degree from Harbin Engineering University, China, in 2012, and the Ph.D. degree from Nanyang Technological University, Singapore, in 2018. She was a Scientist with the Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR), Singapore, until August 2019. Since September 2019, she has been with the Electronic Engineering Department, Fudan University, China, where she is currently a pre-tenure Associate Professor. Her research interests include machine learning, pattern recognition, control and optimization, and complex networks.

...