Ship Segmentation in Aerial Images for Maritime Surveillance

Carlos Pires[1], Alexandre Bernardino[1], and Bruno Damas[1,2]

[1]ISR LISBOA / LARSyS, [2]CINAV

Introduction

Countries with large exclusive economic zones require efficient and cost-effective tools for maritime surveillance. Using Unmanned Aerial Vehicles (UAVs) for automatic ship detection, we release a considerable amount of resources. Ship segmentation in the acquired images is a challenging task, as it is affected by factors like scale, perspective, and illumination variations. Since we are dealing with real-time systems, and UAVs have limited processing capacity, we need to implement an efficient and fast segmentation algorithm.

Contribution

We propose a two stages system to perform ship segmentation. We present a cascade model with detection and segmentation, which makes the ship identification faster. In the first stage, we use a fast detection network to search for possible ship locations regions, then we pass these regions through a segmentation network, thus narrowing the image region where segmentation is performed, which significantly improves the overall image processing time.

Method

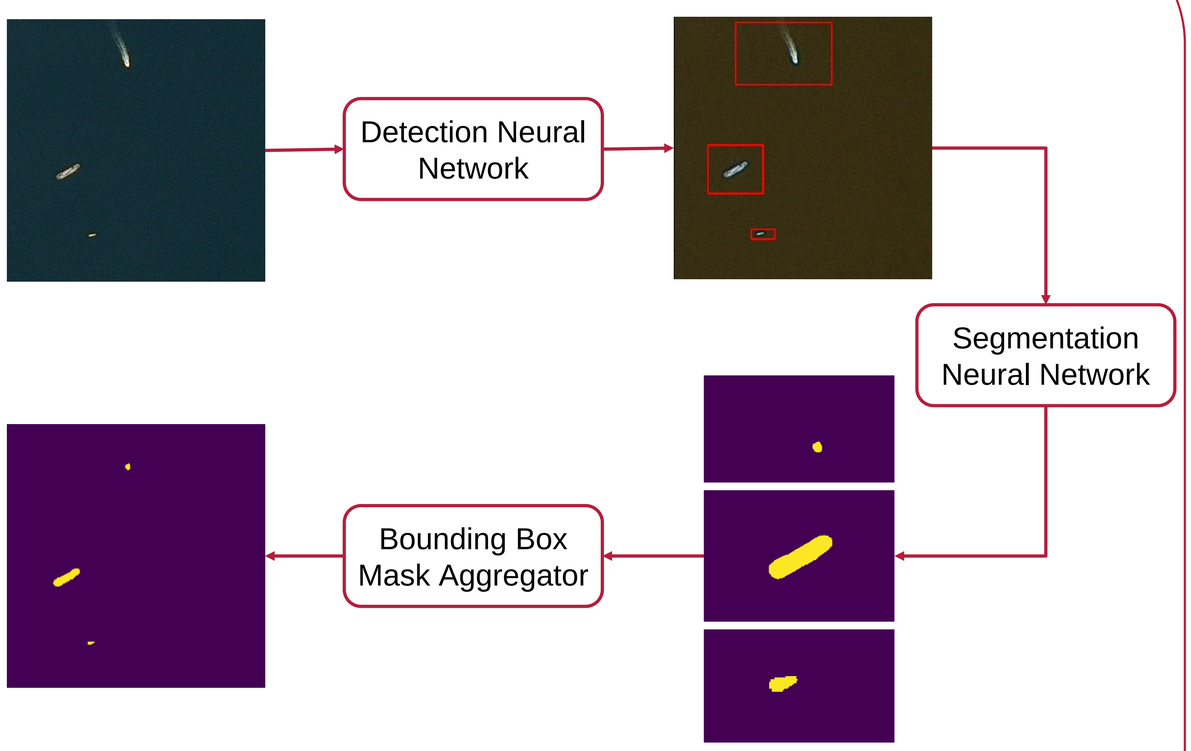

For the detection stage, we implemented the YOLO network [1] with a detection threshold. In this phase, we need a system with high recall, so regions with waves or sun glare are mostly like to be in the region proposals. We used the U-Net to perform semantic segmentation [2]. U-net has an encoder-decoder architecture. So, for the encoder, we can try different architectures and study how it affects performance. The bounding box (BB) mask aggregator module groups the BBs to build the full segmented mask.

Results

To evaluate our method, we tested it with Seagull [3] and Kaggle images [4], also in the Kaggle Airbus ship detection challenge.

A. We tried different thresholds in the detection network.

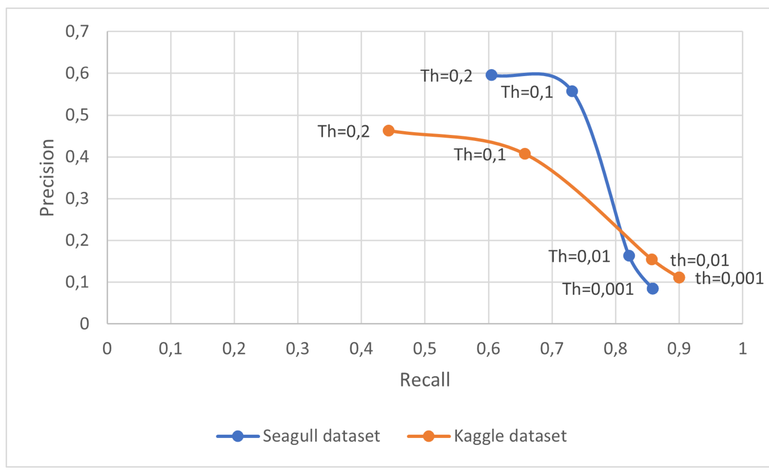

B. We tested the YOLO network on Kaggle and Seagull datasets.

Dataset	Recall	True Positive	False Negative	False Positive
Kaggle	0,9	63	7	504
Seagull	0,86	115	19	1233

C. We tested the U-Net on Kaggle, and Seagull BBs sets.

Dataset	Total Number of BB's	IoU
kaggle	768	0,91
Seagull	494	0,91

D. We compared our method with full image segmentation.

Method	Kaggle Score	Time/image [s]
Full image segmentation	0,71	1,47
Cascade Model	0,82	0,10

Conclusion

This paper presents a contribution to performing fast and accurate maritime ship segmentation. Despite the high number of false positives (FP), the initial stage removes unnecessary parts of the image. Later the segmentation network discards the FPs cases. This cascade approach not only improves the segmentation quality but also greatly decreases the processing time.

Acknowledgements

This work was supported by FCT with the LARSyS – FCT Project UIDB/50009/2020 and project VOAMAIS (02/SAICT/2017/31172).

References

[1] Joseph Redmon *et al.*, Yolov3: An incremental improvement, CoRR, abs/1804.02767, 2018.

[2] Olaf Ronneberger *et al.*, U-net: Convolutional networks for biomedical image segmentation, In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015.

[3] R. Ribeiro *et al.*, A dataset for airborne maritime surveillance environments, In IEEE Transactions on Circuits and Systems for Video Technology, 29(9):2720–2732,2019.

[4] Airbus. Kaggle Airbus ship detection challenge, https://www.Kaggle.com/c/airbus-ship-detection/overview.