

Master Thesis Defense Tilemachos Bontzorlos

Date: Friday, 28.06.2019

Time: 10:30

Room: C119

Title:

An Evaluation of Deep Learning Object Detection Pipelines for Maritime Application Purposes

Abstract:

Maritime is an important sector of modern life and many goods are transported through sea. As a result, there is intense research on Unmanned Surface Vehicle (USV). Autonomous or semi-autonomous ships will decrease the need for personnel on-board, reducing the cost and freeing living space for more cargo. Detection of nearby objects is an essential part of the required autonomy. In more detail, object detection using an optical sensor is a necessary step, since Radar systems can only detect metal objects and at a distance of minimum 1 kilometer. Therefore, an optical sensor can provide information on non-metal maritime objects and objects closer than 1 kilometer. Furthermore, a usual — and expensive — threat in maritime is piracy. Pirates attack vessels using plastic boats, which renders their detection through Radar impossible. An on-board object detection system using optical sensors can provide the crew with early warning of potential threats. Moreover, surveillance of sea borders using USV could benefit of improvements in object detection.

There has been intense research on object detection for maritime purposes. However, in the past years, there has been a breakthrough in object detection through Deep Learning (DL) techniques. These techniques provide accurate and fast detection of objects and have been widely employed in autonomous driving projects. Although there are proposals of object detection using DL in the maritime domain, there is currently limited work on comparison of how different DL architectures perform in a maritime environment.

In this work, different DL architectures using various feature extractors have been compared on a novel maritime dataset, the Singapore Maritime Dataset (SMD). The metrics used for evaluation are the mean Average Precision (mAP), the speed of inference and the maximum usage of the Graphic Processing Unit (GPU) memory.

The results have shown that the Faster R-CNN architecture using Inception v2 for feature extraction has achieved the highest mAP of 76%, although it is not able to run in real-time. On the other hand, SSD architecture has maintained a good ratio of speed and mAP, with up to 40 Frames per Second (FPS) inference speed and up to 65% mAP depending on the feature extractor used. Furthermore, its GPU memory usage has been significantly lower than the other architectures and as low as just 2GB. The results acquired using the SMD have been compared and qualitatively validated with the results of a similar work on the — also novel — SeaShips dataset.
